



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 0 869 439 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.10.1998 Bulletin 1998/41

(51) Int Cl.⁶: G06F 11/14

(21) Application number: 98660029.4

(22) Date of filing: 01.04.1998

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Chin, Howey Q.
San Jose, CA 95124 (US)
• Chan, Kurt
Roseville, CA 95746 (US)

(30) Priority: 01.04.1997 US 825683

(74) Representative: Brax, Matti Juhani
Berggren Oy Ab,
P.O. Box 16
00101 Helsinki (FI)

(71) Applicant: Gadzoox Networks, Inc.
San Jose, Ca. 95119 (US)

(54) Hierarchical storage management from a mirrored file system on a storage network segmented by a bridge

(57) A system for hierarchical data storage management and transparent data backup in a high speed, high volume Fibre Channel Arbitrated Loop environment comprising first (10) and second (26) Fibre Channel Arbitrated Loops, each coupling a Transaction Server (16) and backup HSM server (30) to high speed disk drives (12,14) and mirrored high speed disk drives (32,34) respectively. The two loops are coupled by a Bridge (28) compatible with the Fibre Channel Arbitrated Loop pro-

tol which forwards write transactions directed to the mirrored disk drives (32,34) from the first loop (10) to the second (26) but keeps read transaction from the Transaction Server (16) to the high speed disk drives (12,14) on the first loop (10) isolated from backup and HSM transactions occurring on the second loop (26) between the backup HSM server (30), the mirrored disk drives (32,34) and backup storage devices (38,40) coupled to the backup HSM server (30).




FIG. 1

EP 0 869 439 A1

D scripti n

Background of th Inv ntion

The invention pertains to the field of backup systems in high transaction performance, high data availability networks such as Fibre Channel local area networks in large data storage and processing facilities such as credit card transaction processing centers etc.

In typical high transaction performance, high data availability networks comprising one or more Transaction Servers coupled by a Fibre Channel Arbitrated Loop (hereafter FC-AL or the main loop) to huge banks of on-line storage disk drives called JBODs. Transactions are processed and data is written to the disk drives and read from the disk drives. Typically such read and write transactions keep the main loop busy all the time. However, because of the frequency of failures of disk drives and the uneven frequency of need for read transactions on some data as compared to other data, there is a need in such systems for hierarchical storage management functions such as data aging storage and automatic, nonintrusive backup.

One mirrored system of which the applicants are aware is a mirrored storage system marketed by Vinca. In this system, one Transaction Server is coupled by a SCSI bus to a JBOD, RAID or other hard disk arrangement. This server is connected by a proprietary high speed link to another server. The second server is connected by another SCSI bus to a mirrored JBOD, RAID or other hard disk arrangement. In this arrangement, each time the Transaction Server writes data to its hard drive array, the same data needs to be converted from SCSI format, packetized for transmission over the proprietary high speed link and transmitted to the second server using whatever protocol is in use on the proprietary high speed link. When the second server receives the packets, it must depacketize them and convert the payload data to SCSI format data and then initiate a SCSI write transaction to write the data to the hard disk array coupled to the second server. Thus, the overhead of two SCSI write transactions, packetizing and depacketizing to/from the packet format on the proprietary high speed link and a transmission over the proprietary high speed link must be made to get the data mirrored. The increased overhead increases the server processing loads and thereby increases the cost in CPU time of each transaction. The protocol conversion between SCSI and the protocol on the proprietary high speed link between the servers forces the servers to double as pseudorouters.

Other backup products that are commercially available are software products like the Replica backup software from Stac Electronics in Carlsbad, California. This product can do nothing to solve the problem solved by the invention in speeding up read transactions on a FC-AL since it cannot isolate backup devices from fast main loop devices and there is no mirrored storage. This

means that transactions to the main storage devices cannot be started until the backup transaction from the main storage devices to the backup storage devices are completed.

The requirement for hierarchical storage management creates a conflict in system performance and cost considerations. Hierarchical storage management functions are typically implemented with low performance, low cost devices whereas online transaction processing systems are implemented with high performance, high cost devices. The mixing of hierarchical storage management devices on the same FC-AL with high speed, high performance on-line transaction processing devices results in a significant reduction in overall on-line transaction processing performance. This is because the high performance devices must wait for the low performance hierarchical storage management devices to complete their tasks before the high speed transaction processing devices can continue their work. This is because only one pair of devices can have control of the Fibre Channel Arbitrated Loop at any particular time, so when a hierarchical storage management server has control of the FC-AL to carry out a write transaction to a backup disk drive, no high speed transaction processor can simultaneously be using the FC-AL to do either read or write transactions to the JBOD drives. This becomes an intractable problem as on-line transaction processing systems expand over terabytes of storage and will be becoming worse over time as huge data structures like image, voice and video files are added to these already large data structures.

Therefore, a need has arisen for a way to implement hierarchical storage management functions in such high performance, online transaction processing systems without severely negatively impacting the performance of such systems.

Summary of the Invention

A solution to this problem is provided according to the teachings of the invention by mirroring the data over a Bridge. The Bridge segments the FC-AL into two FC-ALs or two LANs of any other type. One of these LANs carries high speed, high data availability online transaction processing traffic between a Transaction Server and a primary memory, typically comprised of a high performance JBOD or other high speed storage devices. This first LAN or local area network will hereafter be called the Primary Loop or primary FC-AL. The second LAN or FC-AL will be called the Secondary Loop or secondary FC-AL. It carries hierarchical storage management (hereafter HSM) traffic between an HSM backup server and a mirrored memory, typically also a high performance JBOD or other high speed storage device. The Secondary Loop also carries data between the mirrored memory and backup/archival storage devices such as tape drives or WORM drives. The Primary Loop and Secondary Loop referred to herein are preferably

FC-AL, but they may be any other type of local area network such as Ethernet, fast Ethernet, ATM etc.

The way this system works is as follows. For any write transaction from a high speed transaction processor to a high performance JBOD over the Primary Loop, an additional write operation is carried out over the Bridge and the Secondary Loop to the mirrored-storage HSM disk drives. The write transaction to the mirrored disk drive device or devices on the Secondary Loop so as to mirror the data stored on the disk drives in the Primary Loop does not slow down processing on the Primary Loop because of the presence of the Bridge. If the Bridge were not present, a write transaction to a backup storage device such as a mirrored backup disk drive stalls all further transactions with the fast storage devices in the Primary Loop until the transaction with the backup storage device is completed. However, the mirrored disk drives are fast devices so write transactions to them are very fast even though the data must traverse the Bridge. Thus processing of further transactions on the Primary Loop is not appreciably slowed down by a write transaction to a mirrored disk drive on the Secondary Loop. Later, data on the mirrored disk drives in the Secondary Loop is moved by an HSM server on the Secondary Loop to slower backup devices like streaming tape backup drives, WORM drives etc. Since all nodes involved on this secondary backup operation are on the Secondary Loop which is insulated from the Primary Loop by the Bridge, the secondary backup HSM transactions do not tie up the Primary Loop at all.

In contrast, in a prior art FC-AL network without a Bridge, whenever an HSM server started moving data from the mirrored storage to the backup devices such as streaming tape or WORM drives, transactions between the transaction processor and the main disk drives could not occur because the single FC-AL loop to which all nodes were coupled was tied up with HSM transactions to backup devices. With the architecture according to the teachings of the invention, the presence of the Bridge isolates the secondary HSM transactions between the backup devices on the Secondary Loop so that the Primary Loop is not tied up in a loop tenancy and is free to carry out a concurrent transaction between the Transaction Server and the Primary Loop disk drives.

Read operations on the high performance FC-AL segment are unaffected since the Bridge keeps the read operation confined locally to the Primary Loop so no loop tenancy on the Secondary Loop results from a read transaction on the Primary Loop. This provides a major advantage since read operations can outnumber write transactions on the Primary Loop by factors of 9 to 1 in many installations. Thus, these read transactions between the transaction processor and the disk drives on the Primary Loop can proceed at a very fast pace on the Primary Loop and need not be slowed down by slower HSM transactions occurring on the Secondary Loop. That is, HSM transactions on the HSM FC-AL segment

are performed between the backup storage devices, the HSM server and the mirrored disk storage all of which are coupled to the Secondary Loop, and are kept out of the high speed primary FC-AL segment by the Bridge. In this way, 16 to 24 hour or overnight backups no longer have any impact on system operation on the Primary Loop. The only time penalty that is traded off for the acceleration of read transactions is that there is a small latency in the Bridge for each write transaction between the Transaction Server and a destination node on the Primary Loop while the Bridge looks up the destination address in the forwarding table or learns the location of the destination node and decides to keep the traffic on the Primary Loop. This small time penalty is more than offset by the great increase in the read transaction processing rate and the greater security of being able to do data backup and other HSM functions more frequently without slowing down processing on the Primary Loop.

Further, if the Bridge were not present, every primitive and data frame of every transaction on the Primary Loop would have to pass through all the nodes located on the Secondary Loop of the architecture according to the teachings of the invention because they would all be located on the Primary Loop if the Bridge were not present. Thus, each primitive and data frame in FC-AL protocol transmissions would suffer the 6 word delay in each of these nodes even though none of the nodes of the backup machinery is directly involved in the primary transaction.

The architecture and protocol according to the teachings of the invention provide an overall performance boost over nonsegmented, high speed online transaction processing FC-AL networks with integrated HSM components since, in the typical online transaction processing center, read operations dominate write transactions by a large factor according to the information available to the inventors. FC-AL networks in general provide a performance boost as FC-AL networks can be faster than SCSI bus networks (often as much as 5 times faster). SCSI bus networks have a performance curve expressed in I/O transactions per second that rises linearly with increasing numbers of disk drives coupled to the Transaction Server by SCSI buses until a certain level is reached. At that level, performance levels off with increasing numbers of devices and ultimately performance falls off with still larger numbers of devices coupled to the Transaction Server by SCSI buses. The reason for this falloff is that SCSI bus transactions involve a fair amount of overhead processing to set up every I/O transaction. As the number of devices rises, the amount of overhead processing also rises with the number of drives and the number of buses. Ultimately, the overhead of setting up the transaction with the various drives begins to catch up with and overtake the performance increase with increased numbers of drives. It is this phenomenon which causes performance to fall off.

Because Fibre Channel Arbitrated Loops are substantially faster than SCSI buses, there is a movement in the data storage industry to move away from SCSI buses and toward FC-AL networks. In FC-AL networks, more drives can be added than is the case in SCSI bus networks before the performance begins to fall off. This is because of the higher speed of FC-AL networks. Thus, it is easier to attain one terabyte of storage in an FC-AL network without sacrificing performance than in SCSI bus networks.

Further, most Transaction Servers today use PCI buses internally. A maximum limit of 3 SCSI buses can be coupled to a PCI bus. This is because more than three SCSI buses per PCI bus couples too much parasitic capacitance load to the PCI bus thereby slowing it down. Thus, attaining a one terabyte storage system with SCSI buses also requires the use of more Transaction Servers than is the case in a FC-AL network because of the limit on the number of SCSI buses that can be coupled to a PCI bus. The large number of SCSI buses to the JBOD and RAID arrays to reach one terabyte of storage therefore dictates that more Transaction Servers must be used thereby increasing the overall cost and complexity of the system.

Further, since there is no protocol conversion in a network constructed according to the teachings of the invention, the overhead cost in CPU time to get all write data copied to mirrored storage disk drives on the Secondary Loop is much less than for the Vinca prior art where a protocol conversion from the SCSI protocol to the protocol of the proprietary high speed data link between the primary and HSM servers is necessary. Further, the Transaction Server and HSM server in a system according to the teachings of the invention need not have software that can perform a pseudorouter function or the type of protocol conversion needed in the Vinca prior art.

Additional hardware enhancements to support simultaneous writes to the Bridge and the online storage JBODs and enhancements to the Fibre Channel Arbitrated Loop protocol supporting multicast to a Bridge and one or more other nodes can further enhance performance without loss of HSM functionality.

Brief Description of the Drawings

Figure 1 is a block diagram of a system employing the teachings of the invention.

Figure 2 is a diagram showing the structure of one embodiment of a half Bridge according to teachings of the invention.

Figure 3 is a diagram showing how two half Bridges according to the first embodiment of the teachings of the invention can be cross connected internally to make a full Bridge.

Figure 4 is a diagram showing how two half Bridges constructed according to any embodiment of the teachings of the invention may be connected together to link

two Fibre Channel Arbitrated Loops and achieve acceleration over a single FC-AL having the same number of nodes as the two smaller FC-ALs linked by the two half Bridges.

Figure 5 is a diagram showing the internal construction of a half Bridge according to the preferred embodiment of the invention.

Figure 6 is a diagram showing the internal construction of a TX port of a half Bridge according to the preferred embodiment of the invention.

Figure 7 is a diagram showing the internal construction of an RX port of a half Bridge according to the preferred embodiment of the invention.

Figure 8, comprised of Figures 8A through 8G is a flow chart showing processing by the TX and RX ports of a half Bridge according to the preferred embodiment of the invention to carry out bridging, learning and conflicting OPN preemption decisions.

Figure 9 is a block diagram of a full Bridge using the preferred TX and RX port structures of Figures 6 and 7.

Detailed Description of the Preferred Embodiment

Referring to Figure 1, there is shown a block diagram of a system employing the teachings of the invention. A high speed primary transaction processing Fibre Channel Arbitrated Loop (FC-AL) 10 couples a Primary Storage Bank of Hard Disk Drives comprised of a first bank of hard disk drives 12 and a second bank of hard disk drives 14 to a high speed Transaction Server 16. A secondary FC-AL 26 is coupled to the Primary Loop through Bridge 28 and is also coupled to mirrored disk drives 32 and 34 and a backup & HSM server 30. The hard disk drives 12 and 14 on the Primary Loop 10 as well as their mirrored counterparts 32 and 34 on the Secondary Loop 26 and both servers need to be compatible with the Fibre Channel Arbitrated Loop protocol if Fibre Channel Arbitrated Loops are used for the Primary Loop 10 and the Secondary Loop 26. The Seagate Barracuda is one example of an FC-AL compatible disk drive.

The Transaction Server 16 is coupled via a local area network link such as a 10Base-T twisted pair 18 to a network hub 20 such as an Ethernet™ hub. The hub 20 is coupled to a plurality of workstations such as workstations 22 and 24 by individual network drop lines 21 and 23. These workstations carry out processing of transactions which sometimes require that data be read from the banks of hard disk drives 12 and 14 and sometimes require that data be written thereto. These I/O transactions are carried out by the Transaction Server 16 by carrying out read and write transactions in accordance with the Fibre Channel Arbitrated Loop protocol. In some embodiments, the high speed FC-AL 10 may be implemented using an accelerated protocol utilizing the methods and apparatus disclosed in a co-pending U.S. patent application entitled, "ACCELERATED FIBRE CHANNEL HUB AND PROTOCOL", serial number 08/695,290, filed 08/08/96, which is hereby incorporated

ed by reference. In such an embodiment, the FC-AL 10 would be implemented through an accelerated hub of the type described in "Accelerated Fibre Channel Hub And Protocol" patent application referenced above.

The high speed transaction processing FC-AL 10 (hereafter referred to as the high speed loop) is coupled to a secondary hierarchical storage management Fibre Channel Arbitrated Loop 26 by a learning Bridge 28. The learning Bridge 28 may have the design of the Bridge disclosed in co-pending U.S. patent application FIBRE CHANNEL LEARNING BRIDGE, LEARNING HALF BRIDGE, AND PROTOCOL", serial number xxxxxx, filed 1/23/97, which is hereby incorporated by reference and the essence of which is described further below. As mentioned above, the Secondary Loop 26 couples the Bridge to a backup and hierarchical storage management server 30 (hereafter HSM server) and two banks 32 and 34 of mirrored storage hard disk drives. The HSM server 30 is also coupled via a SCSI bus 36 to two online low speed data storage devices 38 and 40 used for storing archival backups and data having infrequent usage or which has not been used recently. In the embodiment shown in Figure 1, low speed data storage device 38 is a CD-based write-once, read-many (WORM) drive and device 40 is a tape backup system.

The way the system of Figure 1 works is by mirroring write transactions in the high speed loop across the Bridge. That is, every time Transaction Server 16 does a write transaction, it arbitrates for the high speed loop, and when it wins control thereof (more than one transaction processing server may be on high speed loop 10), it transmits an open (hereafter OPN) primitive to the particular disk drive in bank 12 or bank 14 to which it wishes to store data. The disk drive on the high speed loop so addressed replies with an RRDY primitive and then the Transaction Server writes a data frame to that disk drive. This process continues with RRDY and data frames being exchanged until either all the data to be stored has been transferred, or one or the other of the Transaction Server 16 or the disk drive transmits a close (CLS) primitive indicating that no more data can be transmitted or accepted during this transaction. This can happen when a buffer in the disk drive becomes full or the Transaction Server for some reason cannot currently complete sending all the data. If the transaction is cut short before all data is transferred, the Transaction Server reopens the connection later and completes the data transfer.

Assuming that the data storage transaction has been completed, the Transaction Server next begins the process of doing a mirrored write transaction of the same data to the mirrored storage hard disk drives 32 and 34. To do this, the Transaction Server 16 generates an OPN primitive addressed to either the backup and HSM server 30 or directly to one of the mirrored storage hard disk drives in banks 32 and 34. In the preferred embodiment, there is mirrored hard disk drive on the Secondary Loop for every hard disk drive on the Primary Loop and the counterpart drives on the primary and Sec-

ondary Loops have identical capacities. Therefore, the hard disk drives on the Secondary Loop have identical content to the hard disk drives on the Primary Loop at substantially all times, and recovery after a disk crash on the Primary Loop can be made by copying the data from the mirrored drives 32 and 34 onto replacement drives on the Primary Loop. When the mirrored drives 32 and 34 have the same capacity as the primary drives, the Transaction Server 16 can address OPN primitives directly to the mirrored drives. In some species within the genus of the invention, the mirrored drives 32 and 34 can have smaller capacity than the primary drives 12 and 14 to save on costs for the system. In these species, the OPN of the mirrored write transaction will be addressed by the Transaction Server 16 to the HSM server 30 which then finds open space on the "mirrored", i.e., backup drives 32 and 34 and forwards the data to the appropriate disk drive on the Secondary Loop. In these embodiments, the HSM server writes the data to an available portion of the backup disk drives 32 and 34 and later archives the data off these backup drives and onto the secondary backup devices 38 and 40 so as to free up space on the backup drives for more writes from the Transaction Server. In all embodiments, the Transaction Server 16, includes some mechanism to try a write to either the HSM server 30 or the backup disk drives 32 and 34 again later if, when a write to either destination is attempted but fails because the Secondary Loop is busy with another local transaction such as copying data from the backup drives to the secondary backup devices 38 or 40. The details of this "try again later" mechanism in the Transaction Server are not critical to the invention, and any way of detecting that the Secondary Loop is busy and scheduling another attempt to write the data to the backup drives later will suffice for practicing the invention.

Hereafter, the discussion assumes a species where the OPN of the mirrored write transaction is directed to the mirrored disk drives directly. When the OPN primitive reaches Bridge 28, its destination address is determined by the Bridge to be on the Secondary Loop, so the OPN primitive is forwarded to the Secondary Loop, unless the Secondary Loop is busy with a local transaction. If there are conflicting OPNs on the primary and Secondary Loops 10 and 26, the Bridge 28 resolves the conflict in the manner described below which may involve sending a preemptive CLS back to the Transaction Server 16. If the transaction receives a CLS before it completes its write to the Secondary Loop, the "try again later" software routine in the Transaction Server is triggered to reschedule another try later.

Assuming no conflicting OPN situation has arisen, and the OPN from the Transaction Server 16 is forwarded onto the Secondary Loop, the following things happen. On the Secondary Loop, the OPN primitive propagates to the destination disk drive. The destination disk drive then replies with an RRDY or CLS primitive which propagates around the Secondary Loop 26 until it reaches

the Bridge 28. When the Bridge 28 sees the RRDY or CLS primitive returned from the Secondary Loop, it concludes that the source is on the Primary Loop and the destination is on the Secondary Loop. Therefore, the Bridge connects the primary and Secondary Loops together as one big loop. The RRDY or CLS primitive thus propagates back to the Transaction Server 16 and the data write transaction is completed in the same fashion as the original write transaction to the disk drive on the Primary Loop was completed. Thus, the same data that was written to the disk drive on the Primary Loop is also written to the mirrored disk drive on the Secondary Loop. After this transaction is completed, the loop tenancy is relinquished, and the Bridge 28 separates the Primary Loop and Secondary Loop again so as to keep all purely local traffic on each loop so two purely local concurrent loop tenancies on the primary and Secondary Loops can coexist.

The backup and HSM server 30 is free to use the Secondary Loop 26 whenever the Secondary Loop is not tied up in a loop tenancy with the Transaction Server 16. The HSM server 30 uses the Secondary Loop to carry out its hierarchical storage management duties without slowing down the Primary Loop by virtue of the presence of Bridge 28. **[Howey & Kurt: the following passage is inconsistent with the "try again later mechanism" previously discussed] Arbitration for the Secondary Loop will always be won by the Transaction Server 16 as between it and the HSM server 30 since the Transaction Server 16 will have a higher priority. Thus, any time both the HSM server and the Transaction Server are arbitrating for control of the Secondary Loop, the Transaction Server 16 will win and be able to complete its transaction before the HSM server 30 can continue any storage management function involving use of the Secondary Loop 26.** Typical HSM transactions will be to move data that has not been used for a predetermined period or which is older than a predetermined age to archival storage. This is typically done by reading the data to be archived from the mirrored storage drives and then doing write transactions to the CD WORM drive 38 or the tape drive 40.

The Bridge 28 works generally as follows. The OPN primitive issued by the Transaction Server has a destination address which designates the address of the particular hard disk drive to be opened. This destination address is part of the data structure of the OPN primitive. The Fibre Channel Arbitrated Loop learning Bridge 28 monitors loop traffic for the destination addresses of any OPN primitives and monitors other primitives that result from the OPN primitive so as to draw conclusions as to the locations of each node in terms of whether they are on the Primary Loop 10 or the Secondary Loop 26. This process is explained more fully below. The locations of the nodes are stored in a forwarding table stored in memory in the Bridge 28. Each destination address in an OPN primitive issued from either the Transaction

Server 16 or the HSM server 30 (or any other node on either loop) is checked against any entry for that destination address in the forwarding table. When the destination address in an OPN received from the Secondary Loop is found in the table and the table indicates that the destination address is located on the Primary Loop 10, the Bridge forwards the OPN primitive to the Primary Loop 10. Likewise, when an OPN primitive arrives from the Primary Loop, and the forwarding table indicates that the node is located on the Secondary Loop, the OPN primitive is forwarded to the Secondary Loop side of the Bridge. After forwarding of an OPN primitive, the destination device sees its destination address in the OPN primitive and replies either with an RRDY or a CLS primitive which causes the Bridge to conclude that the source and destination nodes are on opposite sides of the Bridge and to switch so as to make connections to couple the Primary Loop and Secondary Loop together as one loop. The data transfer process is then completed in accordance with ordinary Fibre Channel protocol rules. If the Bridge 28 determines that the destination address of an OPN is located on the same loop from which it was received, the OPN is not forwarded and the transaction is completed on a purely local basis such that the nodes on the loop having neither the destination node or the source node are bypassed. There follows a more detailed discussion of the construction and operation of the full Bridge 28 and the half Bridges of which it is constructed.

Half Bridge Operation

Referring to Figure 2, there is shown the structure of a half Bridge according to one alternative embodiment of the invention wherein the TX port sets data in the memory during the learning process. Figure 2 shows how a half Bridge can be used to divide a Fibre Channel Arbitrated Loop into two loop segments - a local loop segment 52 and a remote loop segment 54, although half Bridges are never used alone to do this and are always coupled to another half Bridge. The remote loop segment is comprised of two segments: a TX or transmit segment coupling the RX port of the left half Bridge to the TX port of the right half Bridge; and an RX or receive segment coupling the TX port of the left half Bridge to the RX port of the right half Bridge. For purposes of illustrating generally how a half Bridge according to the teachings of the invention, Figure 2 does not show the other half Bridge, but represents it by remote loop segment 54. The half Bridge, when used in conjunction with another half Bridge to form a full Bridge, prevents any local loop traffic from being transmitted into the remote loop segment across the full Bridge as the half Bridges learn the addresses of the nodes on the local loop segment.

The half Bridge has a TX port 56 having internal terminals 1, 2 and 3 which are coupled to a state machine (not shown) which implements part of the switching

rules of the half Bridge. Likewise, the half Bridge has a RX port 58 which also has terminals 1, 2 and 3 which are also coupled to a state machine (not shown) which implements the remaining switching rules of the half Bridge. The details of the internal construction of the TX and RX ports is not critical to the invention. For example, the structure described below for the preferred embodiment can be used with suitable modification of the control signals that flow between the TX and RX ports. Alternatively, other constructions for the TX and RX ports may also be used such as control by a programmable machine or off-the-shelf FC-AL chipsets to implement the learning and switching rules described herein. The TX and RX ports of Figure 2 work generally the same as the TX and RX ports 100 and 102 of Figures 6 and 7 and in terms of switching rules, learning, handling ARB fill words and RRDY primitives, and resolving concurrent OPN situations. The main difference between the embodiment of Figure 2 and the preferred embodiment of Figure 5 is in how the TX and RX ports control and use memory 78. Accordingly, much of the discussion of the specifics of these common operations is deferred until the discussion of Figures 6 and 7. The description of the general operation of the half Bridge of Figure 2 will be useful to understand the basic principles of FC-AL bridging before the detailed discussion of implementation of the TX and RX ports of Figures 6 and 7 swamps the reader in detail.

Terminal 3 of the TX port 56 is coupled to the outbound segment of local loop segment 52 of the loop and terminal 1 is coupled to the inbound segment of the remote loop segment 54 of the FC-AL. Terminal 1 of the TX port is coupled to a local return segment line 60 which is also coupled to terminal 3 of the RX port 58. The local return segment line 60 is used to route primitives and data which originated from a node on the local segment 52 which are directed to another node on the local segment from the RX port directly back to the TX port so that they need not propagate around the remote segment 54. Also, the RX port 58 contains a latch (not shown but shown in Figure 5 as latch 200) which is coupled so as to be able to store the source and destination addresses of OPN primitives received at a Local RX Port, i.e., pin 1 of RX port 58 coupled to the inbound portion of the local loop segment 52. Likewise, the TX port 56 includes a latch 202 (not shown but shown in Figure 5 as latch 202) which is coupled so as to be able to store OPN primitives received at the Remote RX Port, i.e., pin 2 of the TX port 56 coupled to the inbound portion of the remote loop segment 54. The operation of these latches to support resolution of various concurrent OPN scenarios will be explained further below.

The half Bridge 50 is a learning Bridge in that it monitors traffic on the loop segments to which it is coupled and uses the destination addresses of OPN primitives in conjunction with the ensuing RRDY and CLS primitives and the terminal on which they arrive to draw conclusions regarding whether a node having a particular

address is on the local segment or the remote segment. How this is done will be explained in more detail below, but for completeness here, the half Bridge keeps a forwarding table in memory 78 which is consulted by the half Bridge when each new OPN primitive arrives. The half Bridge uses the destination address of the OPN primitive as an index into a forwarding table stored in a memory 78 to obtain data indicating whether the node having the destination address in the OPN primitive is on the local segment 52 or the remote segment 54.

The half Bridge can only short circuit data frames and primitives through the local segment return 60 for the local segment 52 and not for the remote segment 54. Data frames and primitives propagating on the remote segment 54 arrive at TX port 56 and propagate around the local segment and through the RX port 58 as if the half Bridge was not really there. However, primitives and commands that are propagating on the remote segment 54 cannot be short circuited by the half Bridge and must propagate around the local segment 52 even if the destination node is on the remote segment.

The reason for existence of the half Bridge is that it makes local loop segmentation and acceleration possible and is particularly useful in situations where there are two separate concentrations of nodes, each concentration separated from the other by some appreciable distance. In such a situation, symbolized by Figure 4, the first concentration of nodes are coupled by a segment 1 FC-AL designated 62 which are coupled to a first half Bridge 66. The second concentration of nodes are coupled by a segment 2 FC-AL 64 which is also coupled to a second half Bridge 68. This situation requires the use of a long fiber optic link represented by fibers 70 and 72. Fiber 70 couples the RX port of half Bridge 66 to the TX port of half Bridge 68. Fiber 72 couples the RX port of half Bridge 68 to the TX port of half Bridge 66. The two half Bridges in this configuration cooperate to implement a full Bridge which learns the locations of each node on the first and second segment FC-ALs by watching the traffic on the loop. As the half Bridge 66 becomes smarter, more traffic which is between nodes on the first segment 62 gets short circuited in the half Bridge 66 through that Bridge's local return segment. This prevents purely local traffic on loop segment 62 from having to incur the propagation delays inherent in propagating across the fiber optic links 70 and 72 and the latency of each node on the second segment 64 in order to get back to the destination node on the first segment. Likewise, as the half Bridge 68 becomes smarter, more traffic which is between nodes on the second segment 64 gets short circuited in the half Bridge 68 through that Bridge's local return segment thereby accelerating purely local operations on the second segment 64 like half Bridge 66 accelerates purely local operations on the first segment 62.

The way the half Bridge of Figure 2 functions to carry out its learning process and perform switching to ac-

celerate transactions where possible is as follows. Memory 78 stores a forwarding table. The forwarding table is a 1024 x 1 memory which has 1 memory location for each of the 1024 different possible destination addresses that could be generated using a 10 bit AL_PD destination address in an OPN primitive. The reader will note an apparent discrepancy in that the Fibre Channel Arbitrated Loop protocol only permits 126 different destination addresses with destination address numbers 0, F0 (hex), F7 (hex) and F8 (hex) reserved whereas 1024 different destination addresses can be generated using 10 bits. The 126 different possible destination addresses which are actually useable represent 128 addresses possible with a 7 bit address after reserved addresses F0 and 0 are removed. The reserved addresses F7 and F8 are used as a preemption tie breaking method in the event of duplicate half-duplex OPNs. Fibre Channel Arbitrated Loops use 8b/10b encoding and decoding.

In operation, there are several different possibilities for the states entered by the TX and RX Port State Machines and the switching and other processing that will be carried out thereby. The switching and other processing which is actually carried out by the state machines in any situation depends upon the relative locations of the source node and destination node on the local and remote loop segments.

The state transitions, switching rules and other processing carried out by the transmit and RX ports 56 and 58 in Figure 2 for situations other than handling of conflicting OPNs are as follows. Switching rules for handling conflicting OPNs are as given above in the summary of the invention section. The details of the circuitry and/or programming to implement these rules to handle conflicting OPNs and do the switching described herein are not critical to the invention.

The TX port and RX port each initialize such that the TX port makes a 2-3 connection between terminals 2 and 3 therein and the RX port makes a 1-2 connection such that the remote loop segment 54 and the local loop segment 52 are coupled together as one big loop in coordination with the initialization condition that every destination node is assumed to be on the remote loop segment. Memory 78 is cleared to all 0s using the Reset All signal line 500 (1 = local, 0 = remote). The 2-3 connection terminology will be used herein as a shorthand expression meaning making a connection between terminals 2 and 3 and likewise for any other terminal pair. Suppose now the TX port 56 receives an OPN primitive on terminal 2. The TX port of the half Bridge latches the AL_PD (destination address) contained in the OPN primitive in an internal latch (not shown but like latch 202 in Figure 5), places this destination address on address bus 108 and forwards the OPN primitive out onto the local loop segment 52. The TX port 56 then arms the RX port 58 by activating the Arm signal on line 65. This sets the RX port to a certain state that will cause certain processing to occur if a CLS or RRDY primitive arrives at terminal 1 of the RX port from the local loop segment

52. Arming of the RX port is symbolized in Figure 2 by activating a signal on line 65 in Figure 2, although, in the preferred embodiment, the state machines within the TX port 56 and the RX port 58 are the same state machine so all the control signals between the transmit and RX ports shown in Figure 2 would not be necessary as the data conveyed by these control signals would be known by the state machine by virtue of the state it is in and the transitions between states. In alternative embodiments, the TX Port State Machine can be a separate state machine from the RX Port State Machine. Such embodiments typically use the control signals shown in either Figure 2 or Figure 5 flowing between the transmit and RX ports to keep the two different state machines apprised of the conditions encountered by the transmit and RX ports.

If the TX port 56 receives an OPN primitive on terminal 1 from the local segment return 60, it automatically makes a 1-3 connection to forward the OPN primitive out on the local loop segment 52 in case the destination node is there.

If the RX port 58 receives an OPN primitive on terminal 1 from the local loop segment 52, it decides what to do by first deciding whether it has been previously armed. If the RX port 58 has been previously armed, it knows the OPN came from the remote loop segment and its destination node is not on the local loop segment 52 because if it had been local, the destination node would have converted the OPN to an RRDY or CLS for transmission to terminal 1 of the RX port. Thus, the RX port knows that the OPN must be forwarded back to the remote loop segment, so the RX port makes a 1-2 connection to forward the OPN primitive out on the remote loop segment 54. The RX port then writes a logic 0 into the forwarding table address currently latched on the address bus 108 indicating that this particular destination is not on its local segment.

If the RX port 58 in Figure 2 was not previously armed when the OPN primitive arrives on terminal 1, it knows the OPN was generated by one of its local nodes, but it does not know whether the location of the destination node has been previously stored in the forwarding table. In this situation, the RX port latches the destination address of the OPN primitive and places the destination address on the address bus 108. The RX port then reads the memory location mapped to that destination address by activating the Read signal on line 67. Whatever data is stored at that memory location is then output by memory 78 on data bus 112 and causes the RX Port State Machine to enter one of two possible states. If the data returned by the memory is a logic 1, it means that the destination node is local. In this event, the RX port makes a 1-3 connection to transmit the OPN primitive out on the local segment return 60. The RX port also then activates a Local signal on line 71 thereby informing the TX Port State Machine to make a 1-3 connection to keep the traffic local and bypass the delays imposed by the nodes on the remote loop segment. If

the data returned from the forwarding table was a logic 0, the destination node is not local, so the RX port 58 makes a 1-2 connection to forward the OPN primitive out on the remote loop segment 54. The RX port then deactivates the Local signal on line 71 thereby causing the TX port 56 to resume the 2-3 connection of its initialization state.

If the RX port 58 does not receive an OPN primitive on terminal 1 but observes either a CLS or RRDY primitive arriving from the local segment on terminal 1, it knows that the destination node is on the local loop segment 52. Accordingly, the RX port activates the Write One signal on line 73 thereby causing the TX port to activate the Set signal on line 75 so as to set to a logic 1 the memory location mapped to the destination address of the OPN primitive previously received by the TX port from the remote loop segment. Next, the RX port determines whether it has been previously armed. If armed, the RX port knows that an OPN primitive previously received by the TX port from the remote loop segment has been forwarded thereby out onto the local loop segment 52. Therefore, the RX port also knows that the OPN initiator is not local, so it makes a 1-2 connection to forward the RRDY or CLS primitives received at terminal 1 to the source node on the remote loop segment. The RX port 58 then also deactivates the Local signal on line 71 to cause the TX port 56 to make or maintain a 2-3 connection. This connection by the TX port 56 permits any data frames arriving on terminal 2 from the source node to propagate out to the local destination node on the local loop segment 52 so as to allow the data transfer to be completed. In embodiments where the TX port and RX Port State Machines are combined into one machine, the single state machine will only have one state which is assumed in this source remote, destination local situation and will automatically assume the 1-2 connection for the RX port and the 2-3 connection for the TX port.

If the RX port 58 had not been previously armed when it received either an RRDY or CLS primitive at terminal 1, it means both the source node and the destination node are on the local loop segment but the destination node is closer to the RX port than the source node such that the OPN primitive issued by the source node arrived at the destination node before it arrived at the RX port 58. The destination node then replied with either an RRDY or CLS primitive which the RX port sees in an unarmed state. In this situation, RX port 58 makes a 1-3 connection and activates the Local signal on line 71 thereby causing the TX port to make a 1-3 connection. In embodiments where the TX port and RX Port State Machines are combined into one machine, the single state machine will only have one state which is assumed in the source local, destination local situation and will automatically assume the 1-3 connection for the RX port and the 1-3 connection for the TX port. Special rules for handling the fatal embrace case of simultaneous or near simultaneous OPNs arriving at a half Bridge from both

the remote and local loop segments will be described in more detail below in the section heading SIMULTANEOUS OPNs.

5 PREFERRED HALF BRIDGE, TX PORT AND RX PORT STRUCTURES

Referring to Figure 5, there is shown a block diagram of the preferred embodiment for a half Bridge. In the embodiment of Figure 5, the RX port 100 does all the setting and resetting of the memory. This differs from the embodiment of Figure 2 where the TX port sets the forwarding table memory locations to logic 1 and the RX port resets them to logic 0 on system initialization or the reception of a LIP initialization primitive. As is the case for the embodiment of Figure 2, the LIP primitive will also cause the RX port 100 and TX port 102 to assume their default connections, 1-2 and 2-3, respectively. In the embodiment of Figure 5, elements having like reference numerals to elements in Figure 2 are the same and serve the same purpose in the combination.

As in the case of Figure 2, the preferred half Bridge is comprised of a TX port 102, and an RX port 100 coupled by a local segment return 60 and coupled to a memory 78 which stores location data for each FC-AL node address as the location for the node is learned to be local or remote. The half Bridge of Figure 5 includes a comparator 402 which functions to compare source and destination addresses of OPNs to resolve concurrent OPN situations. All of these circuits are coupled together by a plurality of different control signals, the function and activation conditions of each which will become clear during the course of the discussion of the flow charts which describe the switching and learning operations carried out by the TX port and RX port.

Figure 6 is a block diagram of the preferred internal structure of the half Bridge TX port 102. The TX port comprises: a state machine 602 (which can be a programmed microprocessor) which functions to provide the logic which implements the switching and concurrent OPN resolution rules, count RRDYs and controls the switch 608 and insertion/substitution/deletion FIFO circuit 610 to do the operations detailed in the flow charts; an AL_PA latch 202 which functions to store the source and destination addresses of OPNs and to automatically substitute hex F7 or hex F8 into the source address field of any half-duplex OPNs; a remote decode circuit 604 which functions to recognize the primitives that arrive at pin 2 from the remote half Bridge and advise the state machine 602; a local decode circuit 606 which functions to recognize primitives that arrive at pin 1 via the local bypass and advise the state machine 602; a switch 608 which controls routing of primitives and data arriving at either pin 2 from the remote half Bridge or pin 1 via the local bypass onto the local loop segment 52 connected to pin 3; and insertion/substitution/deletion FIFO circuit 610 which functions to block OPNs and trailing RRDYS until concurrent OPN situations can be

resolved, regenerate and forward winning OPNs, generate and forward RRDYs that trailed OPNs, generate and transmit preemptive CLSs, generate fill words such as ARB(0) to arbitrate for control of the local loop and substitute ARB(0) for any incoming fill words during the arbitration process or forward them unchanged when the TX port is not arbitrating, and enter a transparent mode where incoming data frames and primitives are passed therethrough unaltered.

Figure 7 is a block diagram of the preferred structure for an RX port 100. The RX port is comprised of a state machine 612 (which can be a programmed micro-processor) which functions to control memory 78 and a switch 614 and an insertion/substitution FIFO circuit 616 as well as receive and send various control signals to implement the switching, learning and concurrent OPN resolution rules detailed below; an AL_PA latch 200 which latches the source and destination addresses of incoming local OPNs from local loop segment 52 and supplies the AL_PA address via bus 201 to the data input of an Insertion/Substitution/Deletion FIFO circuit when an OPN is generated; a decoder circuit 620 which functions to recognize various primitives arriving on the local loop segment 52 and advise the state machine as well as recognizing when a half-duplex OPN arrives thereby causing the state machine to cause conversion thereof in latch 200 to a pseudo-full-duplex OPN; an insertion/substitution FIFO circuit 616 which serves to block OPNs and trailing RRDYs until possible preemption situations can be resolved, generate and transmit OPNs, preemptive CLSs and RRDYs out to the remote half Bridge via pin 2 under control of the state machine 612 and enter a transparent mode where incoming data frames and primitives are passed therethrough unaltered; and switch 614 which serves to route incoming primitives and data at pin 1 and primitives generated by the insertion and substitution FIFO circuit 616 out on either pin 2 to the remote half Bridge or pin 3 to the local bypass.

Before discussing the detailed operation of the preferred embodiment for the half Bridge, a short discussion of full Bridge operation is in order so as to set the context for the discussion of the half Bridge process flow charts.

Full Bridge Operation

Referring to Figure 3, there is shown a full Bridge formed using two half Bridges according to the embodiment of Figure 2. Figure 9 shows a full Bridge formed using two half Bridges according to the embodiment of Figure 5. This full Bridge is comprised of two half Bridges which have the structure of the half Bridge of Figure 5 and function the way described in the flow charts of Figures 8A through 8G. The full Bridge of Figure 3 is two half Bridges connected together to filter traffic from two loop segments 52A and 52B and allow two concurrent loop tenancies, i.e., two local transactions may be occurring simultaneously.

The two half Bridges 149 and 151 shown in the embodiment of Figure 3 have identical structures and work identically to the above description of the half Bridge embodiment of Figure 2.

In the Bridges of either Figure 3 or Figure 9, traffic arriving on terminal 2 of the TX port of either half Bridge is minimized. In other words, after the learning process, only OPNs destined to nodes not on a local segment connected to one of the half Bridges get forwarded to the other half Bridge. For example, in the embodiment of Figure 9, if the source node is node 153 and the destination node is node 155, all data frames traveling between these two nodes are shunted across local return segment 60B and never reaches left half Bridge 161. Likewise, if the source node is node 104 and the destination node is node 106, all data frames traveling between these nodes are shunted across local return segment 60A and never reach right half Bridge 163. An advantage of this structure is that concurrent conversations can simultaneously occur between nodes 153 and 155 as one conversation and nodes 106 and 104 as another concurrent conversation. However, if the source node is 153 and the destination node is 106, the OPN primitive from node 153 will be forwarded on line 54B from terminal 2 of RX port 100B to terminal 2 of TX port 102A in accordance with the switching rules defined above. After TX port 102A wins arbitration with ARB(0), the OPN propagates to destination node 106 which responds with an RRDY or CLS primitive. The response primitive arrives at terminal 1 of RX port 100A and is forwarded via a 1-2 connection and line 54A from terminal 2 of RX port 100A to terminal 2 of TX port 102B where it propagates to the source node 153.

Similarly, as shown in Figure 4, two half-Bridges can be used to connect two groups of nodes that are physically separated by long distances. The two half Bridges 66 and 68 can be of the design of Figure 2 or the design of Figure 5 or any other design within the teachings of the invention. The two half Bridges are connected together by long fiber-optic segments 70 and 72. Since the speed of light through a fiber optic cable is roughly 5 ns/meter, a 200 meter full-duplex link line fibers 70 and 72 between two loop segments results in $5 \times 9 \times (2 \times 200) = 2$ microseconds of additional latency over a short link, which is the equivalent of 9 nodes on the link (40 bits per FC word, 6 FC words per node, 941 picoseconds per bit at 1.0625Gb/sec). On a conventional loop, all transfers would incur this additional 9 node equivalent delay. With additional 200m fiber optic links, additional delay is incurred. For example, with two 200m links the additional delay would be equivalent to 18 nodes, with three links 27 nodes, etc. On a Bridged segment such as is shown in Figure 4, only the traffic flowing across the Bridge would incur these delays.

The operation of the half Bridge embodiment shown in Figure 5 and the TX and RX ports shown in Figures 6 and 7 will be described by reference to the flow charts of Figure 8 comprised of Figures 8A through 8G.

Concurrent half-duplex OPNs present a difficult concurrent OPN resolution problem since the normal preemption rules described below do not work since a half-duplex OPN only has two destination addresses and no source address so an ambiguity is created as to whether they are the same OPN (same source and destination address) or not. This is resolved as follows. First, a rule is adopted that no half duplex OPN can be forwarded from one half Bridge to another. Instead, any half duplex OPN received at the RX port is recognized and automatically converted to a pseudo-full-duplex OPN when it is stored in latch 200 as follows. In a half-duplex OPN where no source address is provided, the half Bridge RX port decode circuit 620 will recognize the arrival of a half duplex OPN at pin 1 and activate the Half Duplex control signal to alert state machine 612 of the event. The state machine 612 then asserts RX Convert to cause latch circuit 200 to insert either hex F7 or hex F8 for the source address in the AL_PA latch 200 in place of the destination address that was originally there. The choice between F7 and F8 depends upon how the Bridge was configured upon installation. This permits easy preemption resolution using the full duplex rules described herein in the event of concurrent OPNs with the same AL_PD. The OPNs are then compared by the preemption process described below and one or the other of them is preempted. One of the full or pseudo-full-duplex OPNs will survive and will be forwarded by the RX port of the half Bridge which received it to the other half Bridge or over the local bypass path. When a winning OPN is forwarded by the RX port to the remote half Bridge, if it was originally a half-duplex OPN, it is not converted back to its original half duplex form. When a winning OPN is forwarded by the RX port over the local bypass, it is converted back to its original half-duplex form by assertion of the RX Convert Back signal by state machine 612. This causes the Insertion/Substitution circuit 616 to substitute the AL_PD of the OPN for its F7 or F8 in the source address field. When a winning OPN is forwarded onto the local loop by the TX port, if it was originally a half-duplex OPN, it is converted back to its original half-duplex form by assertion of a Convert Back signal by TX Port State Machine 602. This causes the Insertion/Substitution/Deletion FIFO circuit 610 to substitute the AL_PD of an OPN for the F7 or F8 source address thereby converting it back to its original form.

SWITCHING RULES AND STATES FOR THE TRANSMIT AND RX PORTS OF THE PREFERRED EMBODIMENT OF THE HALF BRIDGE OF FIGURE 5

There is given below the port state machine state transition and switching rules implemented by the state machines in the RX and TX ports, 100 and 102, respectively, in the embodiment of Figure 5. The rules are given in the form of flow charts. The section of this specification entitled RULES FOR HANDLING OF ARB PRIMITIVES BY BOTH HALF BRIDGE AND FULL BRIDGE

contains arbitration primitive handling rules for both the RX ports and the TX ports of the embodiments of Figures 2 and 5, and is incorporated by reference into the following port state transition and switching rules.

RULES FOR HANDLING OF ARB PRIMITIVES BY BOTH HALF BRIDGE AND FULL BRIDGE

ARB primitives are output by any node or TX port or RX port desiring to obtain control of a loop segment prior to transmission. The ARB primitive includes a priority designator therein which indicates the priority of the node or TX port or RX port that issued it. Each node or TX port or RX port which receives an ARB primitive fill word examines the priority thereof. If that node or TX port or RX port desires to transmit, the ARB priority designator of the incoming ARB is changed to the priority designator of the node or TX port or RX port which desired to transmit if the priority thereof is higher than the priority of the incoming ARB primitive and the ARB is forwarded. If the priority of the incoming ARB is higher than the priority of the node, TX port or RX port which desires to transmit, the incoming ARB is forwarded without change. If a TX port, RX port or node sees its own ARB comes back to it, it knows it has won the arbitration and has permission to transmit. After a node or a TX port or an RX port sees its own ARB come back to it, it thereafter "swallows" all incoming ARBs, regardless of their priority, by substituting the highest priority address, F0, for the priority designator of all incoming ARBs until the node, TX port or RX port has relinquished the loop. At that time, the ARB handling rules defined above for the time before the node saw its own ARB come back are resumed.

In the case where the TX port 102 receives a remote OPN on pin 2, it begins to arbitrate for control of the local loop segment. It can happen that a node on the local loop segment decides that it needs to transmit. That local node will also begin to arbitrate for the local loop segment. If the local node's ARB comes back to it before the remote OPN arrives, the TX port's ARBs will all be swallowed, and the TX port will block the remote OPN from transmission onto the local loop and latch the source and destination addresses thereof. Meanwhile, the local node will generate an OPN which arrives at the Local RX Port, pin 1 of the RX port 100. This will trigger the preemption rules above to decide which OPN to preempt. If the local node begins to arbitrate for the local loop segment after the TX port has arbitrated for and won control of the local loop, no preemption situation arises. This is because the local node's ARBs will be swallowed by the TX port so it will never gain control of the local loop segment and transmit a local OPN to the RX port that would cause a concurrent OPN situation to arise.

On a half Bridge, the ARBs from the local nodes arriving at the Local RX Port, all remain local, i. e., are transmitted out pin 3 of the RX port to pin 1 of the TX

port. The TX port then just forwards the ARB primitiv out on pin 3, the Local TX Port. On a full Bridge, the same thing happens. Specifically, the ARB primitives arriving at the Local RX Port of the RX port are not forwarded to the other half Bridge but are sent only to the TX port of the half Bridge which received the ARB primitive. Consequently, on a full Bridge, no ARBs flow between the local and remote loop segments coupled to the two half Bridges. Instead, ARBs arrive at the Local RX Port and are forwarded to terminal 1 of the local TX port without ever being transmitted on terminal 2 of the local RX port (see Figure 3) to the remote loop segment.

Half Bridges cannot stand alone. They must be connected to another half Bridge either locally to make a full Bridge or remotely as in the configuration of Figure 4.

Referring to Figure 8, which is comprised of multiple pages of flow charts labelled Figure 8A, 8B etc., there is shown a flow chart graphically illustrating the half-to-pseudo-full-duplex conversion processing, learning, preemption and switching rules which are followed concurrently by the TX port 102 and the RX port 100 in each half Bridge of a full Bridge. The reader should read the flowcharts of Figure 8A et seq. in conjunction with study of Figures 5, 6 and 7 for a full understanding of the functions and relationships between the circuit elements and control signal activations. Below there are reiterated 6 possible scenarios some of which require preemption and some of which do not. Processing of the various cases on the flowcharts of Figures 8A et seq. is indicated by labels at appropriate branches of the flowchart.

Case 1: a local OPN at the left half Bridge is forwarded to the right half Bridge followed by receipt at the left half Bridge of a different remote OPN. In this case, conflicting OPNs make preemption of one of them necessary. The preemption decision is based upon a comparison of the addresses of the conflicting OPNs. Half duplex OPNs are converted to pseudo-full duplex OPNs before the preemption address comparison is done. If the address comparison indicates that the local OPN of the left half Bridge is higher priority than the remote OPN received at the left half Bridge, the remote OPN is discarded by the left Bridge and the right half Bridge sends out a preemptive CLS to close the source node which generated the lower priority remote OPN. If the remote OPN at the left half Bridge is higher priority, the left half Bridge transmits a preemptive CLS on its local loop to close the node which generated the lower priority local OPN. The left half Bridge then arbitrates for control of the left half Bridge local loop using ARB(0). Any ARB fill words coming in from the right half Bridge are converted to ARB(0) by the left half Bridge during this process. Arbitration is won by the left half Bridge when it sees ARB(0) return to it. When arbitration is won by the left half Bridge, the remote OPN is transmitted by the left half Bridge out onto the left half Bridge local loop. The right half Bridge, which independently does its own comparison, concludes that the lower priority OPN forwarded to it by the left half Bridge must be discarded.

Cas 2: a local OPN for the left half Bridge has been forwarded across the local bypass to the left half Bridge TX port when, later, a different remote OPN arrives at the TX port of the left half Bridge. This situation is resolved by an automatic preemption of the remote OPN by the left half Bridge sending out a preemptive CLS to the right half Bridge since the local loop is busy.

Case 3: a remote OPN is received at the left half Bridge simultaneously with receipt at the left half Bridge of a local OPN. This situation is resolved by examining the addresses of the two OPNs and sending a preemptive CLS to close the source node of the lower priority OPN, the preemptive CLS being transmitted by the half Bridge coupled to the local loop coupled to the source node of the lower priority OPN. Following the preemption, if the local OPN is higher priority, the memory is accessed to determine whether the destination address of the winning local OPN is local or remote, and the local OPN is forwarded via the appropriate path to its destination.

Case 4: a remote OPN is received by the left half Bridge followed by receipt of a different local OPN. When the remote OPN is received, the TX port starts to arbitrate for the local loop. If a local OPN is received before arbitration is won by the TX port, the remote OPN is too late because the local loop is considered to be busy as soon as arbitration is won by the source node which generated the local OPN. Therefore, the remote OPN is preempted by transmission of a preemptive CLS from the left half Bridge to the right half Bridge, and the left half Bridge then discards the remote OPN.

Case 5 (no preemption necessary): a local OPN is received at the left half Bridge and is forwarded to the right half Bridge, whereupon it is returned to the left half Bridge. When a local OPN is returned from the right half Bridge, the left half Bridge must identify the fact that the remote OPN just received is the same OPN as the local OPN previously forwarded to the right half Bridge. This is done by comparing addresses. When case 5 arises, the returning remote OPN is forwarded transparently onto the local loop without arbitration and all subsequent traffic is forwarded by the left half Bridge transparently onto the local loop until the loop tenancy is done, and a new local or remote OPN is detected.

Case 6 (no preemption necessary): a remote OPN is received, forwarded locally and returns. Upon detection of the identical OPN, the arbitration process started by the left half Bridge TX port is stopped, and the left half Bridge goes into a transparent mode where all fill words, data and primitives coming from the right half Bridge are passed transparently through the local loop and forwarded back to the right half Bridge where they either reach a destination node or are forwarded to the source node. If the destination node is present on the local loop of the right half Bridge, it responds with an RRDY or CLS which is forwarded to the source node on the same local loop. If an RRDY was sent, the source node responds with a frame of data which is then forwarded

warded to the left half Bridge and transparently passed therethrough, through the left half Bridge local loop and back to the right half Bridge where it is transparently forwarded to the destination node. This process is continued until the loop tenancy is completed. The right half Bridge learns the location of the destination node, so the next time a right half Bridge tries to open the same destination node, the local bypass path will be used and the left half Bridge will be omitted from the loop tenancy.

Processing starts in Figure 8A with block 101 wherein the system is initialized by carrying out the following operations to initialize. When the loop initializes (any Loop Initialization Primitive or LIP detected by the TX or RX ports of the half Bridge), the half Bridges connect both local loop segments together as one large loop. Accordingly, TX port 102 makes 2-3 connection, and RX port 100 makes 1-2 connection. This is done by TX Port State Machine 602 setting the state of a Switch Control signal on line 620 so as to set switch 608 to connect pins 2 and 3, and similarly for the Switch Control signal on line 622 generated by RX Port State Machine 612. Also, the memory 78 is cleared to all logic 0s indicating there are no local ports yet and all traffic is to flow through both half Bridges. This is done by the RX Port State Machine 612 asserting the Clear All signal on line 120.

After initialization, test 103 is performed to determine if the TX port has received a remote OPN primitive at the Remote RX Port, terminal 2. If not, processing proceeds to step 818 to determine if the RX port has received a local OPN on pin 1. If not processing proceeds back to Start in block 99.

If the RX port has received a local OPN, processing proceeds to step 822 wherein the Decode circuit 620 in Figure 7 activates an OPN signal on line 654 to inform the RX Port State Machine that an OPN has arrived. The Decode circuit 620 also activates a TX Arm signal on line 644 to tell the TX port that it has received an OPN. The Decode circuit 620 also activates a Latch [0:1] signal on line 656 causing the AL_PA addresses of the local OPN to be stored in latch 200. In response to the activation of the OPN signal on line 654, the RX Port State Machine activates a Del OPN signal on line 727 which causes the Insertion/Substitution/Deletion FIFO 616 to delete the first 20 bits of the header and OPN from the FIFO pipeline thereby blocking the local OPN from being forwarded to the remote half Bridge or on the local segment return 60 until such time as the state machine allows it to be forwarded. If any following RRDYs were received, the Decode Circuit 620 in Figure 7 also activates the RRDY signal on line 730 once for each received RRDY. This causes an RRDY counter in RX Port State Machine 612 to be incremented once for each received RRDY. It also causes the state machine 612 to activate a DEL RRDY signal on line 858 in Figure 7 once for each received RRDY. This causes the FIFO circuit 616 to delete the RRDYs from the FIFO pipeline to block their transmission. Both the OPN and RRDYs can be

reconstructed later for forwarding by the state machine by activation of the Insert OPN and Insert RRDY signals on lines 774 and 786, respectively. The processes of storing the AL_PA and blocking the OPN and any following RRDYs from being forwarded caused by activation of the various signals in step 822 are symbolized by step 824. The RX Port State Machine also activates an Enable Compare signal on line 638 to compare the AL_PA address of the local OPN with the Default AL_PA Address stored in latch 202 of the TX port (the default AL_PA will be stored in both latches 200 and 202 whenever a CLR is received by either RX or TX port).

Step 822 also represents the process of detecting if the incoming local OPN was half-duplex and converting it to pseudo-full-duplex if it was. This process is carried out by the RX port if the Decode circuit 620 activates the Half Duplex signal on line 658 when the OPN signal on line 654 is activated. This causes the RX Port State Machine to activate an RX Convert signal on line 660 in Figure 7 to cause the AL_PA latch circuitry 200 to replace the source address with either hex F7 or hex F8 depending upon the configuration data for the half Bridge.

If a remote OPN arrives later after the local OPN, the address comparison and preemption processing is described below in the description of the processing for preemption cases 1, 2 and 3. The purpose of steps 818 and 822 et seq. is to simply describe the processing which follows when a local OPN arrives alone at the RX port.

Continuing with the discussion of step 822 in Figure 8A, if the local OPN was half-duplex, the Decode circuit 620 also activates a Half Duplex signal on line 658 in Figure 7. This causes the RX Port State Machine to activate an RX Convert signal on line 660 to cause the source address stored in latch circuit 200 to be converted to either hex F7 or hex F8 depending upon the configuration data of the half Bridge.

Further, if the local OPN was trailed with any RRDYs, the Decode circuit 620 activates an RRDY signal for each one. Each time RRDY is activated, an RRDY counter in the RX Port State Machine is incremented.

Finally, after step 824, processing proceeds along path 820 to step 816 in Figure 18E.

Returning to the consideration of step 103 on Figure 8A, the arrival of a remote OPN is detected by the Remote Decode circuit 604 which is coupled by line 630 to pin 2 in Figure 6. As soon as the TX port has received a remote OPN, the Remote Decode circuit 604 activates the OPN control signal on line 632. This tells the state machine that a remote OPN has arrived. When this happens, step 111 is performed to arm the RX port to aid in the learning process, as will be explained below. The TX Port State Machine arms the RX port by activating the RX Arm signal on line 634. Step 111 also represents the process whereby the Remote Decode circuit 604 activates the OPN signal on line 632 which causes the TX

Port State Machine to activate the Del OPN signal on line 691 in Figure 6. This causes the blocking of further transmission of the OPN as symbolized by block 105 by causing the Insertion/Substitution/Deletion FIFO circuit 610 to remove the OPN from the FIFO pipeline. The Remote Decode circuit 604 also activates the Latch [0:1] signal on line 636 which causes the AL_PA latch 202 to latch the AL_PA address of the OPN just received as symbolized by block 105. If the OPN was followed by any RRDYs, the Remote Decode circuit 604 activates the RRDY signal on line 696 for each one. Each such activation causes the TX Port State Machine 602 to increment an RRDY count, and to activate a Del RRDY signal on line 856 for each activation of the RRDY signal on line 696. This removes the RRDYs from the FIFO pipeline to block their transmission. These RRDYs can be later regenerated and transmitted by activation of the Insert RRDY signal on line 686 if the OPN just received wins the AL_PA comparison. This process of blocking the OPN any following RRDYs from transmission through the FIFO is performed by each TX and RX port in both the left half and right half Bridges whenever an OPN is received alone or in combination with following RRDYs and should be understood as the process performed whenever blocking of further transmission of OPNs and RRDYs is required by the flow charts of Figure 8A et seq.

Also, if the received OPN was pseudo-full-duplex, the Pseudo Full Duplex signal on line 710 in Figure 6 is activated by the Remote Decode circuit 604. This informs the state machine 602 to activate the Convert Back signal on line 712 if the OPN is forwarded onto the local loop. This causes the Insertion/Substitution/Deletion FIFO to substitute a copy the destination address for the hex F7 of F8 source address to convert the pseudo-full-duplex OPN back to half duplex before forwarding out onto the local loop segment.

Activation of Enable Compare on line 638 by Remote Decode circuit 604 causes an address comparison of the AL_PA of the remote OPN from latch 202 with whatever AL_PA address is latched in latch 200 of the RX port. If there is no conflicting OPN, latch 200 will store an AL_PA which will always lose against the AL_PA of the remote OPN. Note that in the steps of the process symbolized by Figure 8A et seq., any step that requires activation of Enable Compare enables this signal only if it is not already enabled.

After arming the RX port, step 105 is performed to latch the source and destination addresses of the remote OPN. Step 105 is accomplished as a result of the activation by the Remote Decode circuit 604 of the Latch [0:1] control signal on line 636 successively when the destination and source addresses arrive at pin 2. This signal is activated once as the destination address arrives to latch the 10 bit AL_PD destination address in latch 202 and then is activated again as the AL_PS address arrives to latch the 10 bit source address in latch 202. Step 105 also blocks transmission of the remote

OPN until its destination can be determined and it is clear there are no conflicting OPNs. This blocking of the remote OPN is accomplished automatically as the Remote Decode circuit 604 in Figure 6 recognizes the OPN and activates an OPN signal on line 632. As noted above, this causes TX Port State Machine 602 to activate a Del OPN signal on line 691 which causes the Insertion/Substitution/Deletion FIFO circuit 610 to strip off the 20 bits that comprise the header and OPN primitive. The OPN can be regenerated later if it needs to be forwarded onto the local loop by asserting an Insert OPN signal on line 694.

Test 103 essentially can be a continuous sampling of the OPN signal on line 632, or it can simply be a change of state by TX Port State Machine 602 when the OPN signal on line 632 changes states.

If test 103 determines that a remote OPN has been received, it may mean that the source node is somewhere on the remote loop segment 54. However, it may also mean that the source node and destination node are both on the local loop segment 52 but the source node is closer to the RX port 100 than the destination node and the location of the destination node is as yet unlearned resulting in the OPN being forwarded by the RX port 100 out onto the remote loop segment 54 whereupon it eventually returns to the Remote RX Port, terminal 2 of the TX port via remote loop segment 167. This is the reason the RX port is armed in step 111. By arming the RX port, the half Bridge can learn the location of the destination node of the remote OPN by watching the traffic that results at the local RX port as a result of receipt of the remote OPN. For example, if the source node was node 104 in Figure 5 and the destination node was node 106, source node 104 could transmit an OPN primitive to terminal 1 of the RX port 100 which would then forward the OPN primitive by making a 1-2 connection if it had not yet learned that node 106 was on its local loop segment 52. This OPN would propagate around the remote loop segment and return to terminal 2 of the TX port. The TX port 102 would react to the arrival of the OPN primitive by storing the OPN addresses in latch 202, and comparing the address fields of the remote OPN latched in latch 202 to the address fields of the latched local OPN stored in latch 200 in the RX Port. This comparison is done by activation of the Enable Compare signal on line 638 in Figure 5 and transmission of the AL_PA address fields of the OPN latched in RX port via a 20 bit data bus 108 and the AL_PA address field of the OPN latch in the TX port 102 via a 20 bit bus 109 to a comparator 402. If the remote OPN was the same as the local OPN, the comparator 402 would activate the Equal signal on line 640 which would cause both the TX and RX ports to go into transparent mode wherein all primitives and data are passed therethrough without change. Since any local half-duplex OPNs (characterized by the third and fourth address characters being identical) are converted to pseudo full-duplex OPNs in the RX port latch prior to transmission to the

other half Bridge, if the above situation arose where a local half-duplex OPN was forwarded to the remote half Bridge and came back therefrom, the comparison would occur on the pseudo full-duplex OPNs and the Equal signal would be activated causing transparent mode to be entered. The conversion of any local half-duplex OPNs to pseudo full-duplex OPNs is done by setting the AL_PS (source address) field to F7 or F8, depending on whether the half Bridge was configured to be the high priority or low priority half Bridge at the time of initial configuration of the system. This conversion happens either before the address fields are latched in the internal OPN latches in the TX or RX ports or after the latching process and before any process of shipping the OPN on bus 108 from the RX port to the comparator 402 and before any comparison process. Note that if both the local OPN received by the RX port and the remote OPN received by the TX port of the same half Bridge were pseudo-full-duplex, the source address of each of the local OPN and the remote OPN would both have been changed to F7 or F8, depending upon which way the RX port of the half Bridge was configured. In the normal case however when a remote OPN arrives which is pseudo full-duplex, the F7 or F8 source address is stripped and the OPN is converted back to half duplex before transmission out on the local loop. If a remote pseudo full-duplex OPN is involved in a preemption comparison, if the winning OPN is the remote pseudo-full-duplex OPN, it is converted back to half-duplex before being forwarded to the local loop segment. If the winning OPN is either a local full duplex OPN or a local pseudo full-duplex OPN, no change is made to the full-duplex or pseudo full-duplex OPN prior to transmission to the remote half Bridge.

After latching the source and destination addresses in step 105, processing proceeds via path 642 to test 652 on Figure 8B. That test checks the TX Arm signal on line 644 in Figure 5 to determine if it has been recently activated. The TX Arm signal is set active by the RX port whenever decoder 620 in Figure 7 detects a local OPN at pin 1. The TX Port State Machine notes this fact and sets an internal flag or enters an internal state that can be checked when step 662 is reached to determine which preemptive processing state is necessary.

When a local OPN has been received at pin 1 in Figure 7, Decoder 20 activates the OPN signal on line 654 which causes the RX Port State Machine to activate the TX Arm signal for the one clock cycle needed to access memory and then resets TX Arm on the next clock cycle. Activation of OPN on line 654 also causes the RX Port State Machine to activate the Latch [0:1] signal on line 656 which causes the addresses of the local OPN to be stored in latch 200. If the local OPN was half-duplex, Decoder 620 detects this and activates the Half Duplex signal on line 658. This causes the RX Port State Machine to activate the RX Convert signal on line 660 which causes the AL_PA latch circuit 200 to convert the source address of the local OPN to hex F7 or F8 depending upon the RX port's configuration data.

Activation of TX Arm is done so that if a concurrent remote OPN arrives during the time the local OPN is being dealt with or just after it is forwarded, the conflict can be resolved. The TX Arm signal is activated as soon as the local OPN is detected, and remains active for the single clock cycle during which the memory 78 is checked to determine from the local OPN's destination address where to send the local OPN. During the clock cycle when the local OPN was received, its addresses are latched, its destination address is used to access memory, and, if it is half-duplex, its source address is converted to hex F7 or F8 in the address latch 200. The TX Arm signal is kept active true for the one clock cycle during which the memory access occurs, and is automatically reset on the next clock cycle.

If test 652 in Figure 8B determines that TX Arm is active when the remote OPN arrived, it means a possible conflicting OPN situation has arisen. Cases 1-4 in the preemption rule processing detailed above in the summary of the invention can result depending upon where, if anywhere, the local OPN has been sent. To determine which preemption case processing is necessary, step 662 is performed to determine where the local OPN was sent or if simultaneous local and remote OPNs have been detected. This is done by determining if TX Arm is still active or is false. If TX Arm is false, but has been recently activated, step 662 checks the status of the RX Switch Pos. signal on line 664 in Figure 5. This signal is set by the RX Port State Machine to a logic state which corresponds with the position of switch 614 when the TX Arm signal is deactivated after the memory access has been completed and the switch moved to the position necessary to forward the local OPN to its destination.

Case 2 arises if step 662 determines that the local OPN has been previously forwarded on the local bypass to the TX port prior to the time the remote OPN arrives. The TX port determines where the local OPN was sent on the local bypass by determining if the TX Arm signal is still true, and, if not, by checking the status of an RX Switch Pos. signal on line 664 to see if switch 614 in Figure 7 is set to the 1-3 position. When TX Arm is false, it means that the memory access has been completed and the RX Switch Pos. Signal will indicate whether the local OPN was sent on the local bypass or forwarded to the remote half Bridge.

When a case 2 preemption situation arises, step 666 is performed to carry out an automatic first-come, first-served preemption. Case 2 involves the local loop being busy at the time the remote OPN arrives. This will be the case when step 662 discovers that the local OPN has been previously forwarded on the local bypass when the remote OPN arrives. Step 666 carries out this automatic preemption by the TX Port State Machine activating the Auto Preemption signal on line 646 in Figure 5. This causes the RX Port State Machine to set the Switch Control signal on line 622 in Figure 7 to set the switch 614 to a 1-2 connection and to activate the Insert

CLS signal on line 650 in Figure 7. This causes the Insertion/Substitution FIFO circuit 616 to generate a CLS primitive and send it out to the remote half Bridge where it is forwarded to the source node which generated the remote OPN thereby closing it.

If step 662 determines that the local OPN just received has been previously forwarded to the remote half Bridge at the time the remote OPN arrived, a case 1 preemption situation has arisen. In this case, step 668 is performed wherein the TX Port State Machine activates Enable Compare signal on line 638 in Figure 5. This causes the Comparator 402 to compare the AL_PA addresses of the local and remote OPNs latched in latches 200 and 202 and activate either the Preempt Local Source signal on line 670 or the Preempt Remote Source on line 672.

Next, step 674 is performed to determine which of these two result signals from the comparator have been activated. The two half Bridges that together comprise the full Bridge each simultaneously perform the processes depicted in the flow chart of Figures 8A et seq. However, the flow charts have been written from the viewpoint that the process depicted is occurring in the left half Bridge as the local half Bridge such that the right half Bridge may be referred to as the remote half Bridge. That being said, path 676 out of step 674 represents the branch taken if the remote OPN at pin 2 of the left half Bridge TX port won the address priority comparison. Path 678 represents the branch taken if the left half Bridge local OPN won.

Referring to Figure 8C, the processing along path 678 will be described. When the left half Bridge local OPN won and has already been forwarded to the right half Bridge (the reader is also referred to Figure 3 for context), the right half Bridge must close the source node on its local loop segment before transmitting the local OPN received from the left half Bridge onto its local loop. This is done by step 688 in Figure 8C. In this step, the right half Bridge first blocks the remote OPN received at pin 2 (the forwarded local OPN received from the left half Bridge). This blocking is done when the Remote Decode circuit 604 (see Figure 6) in the right half Bridge (identical circuits to circuits in the left half Bridge are referred to by the same reference numbers) detects a remote OPN at pin 2 and activates the OPN signal on line 632. This causes the Latch [0:1] signal on line 636 to be activated which causes AL_PA latch to latch the source and destination addresses. Activation of the OPN signal on line 632 also causes the TX Port State Machine to follow the same processing that the left half Bridge TX port does when it receives an OPN, said processing being described by the flow chart of Figures 8A et seq. In this case, since step 652 performed by the right half Bridge would find that TX Arm had been recently activated when the right half Bridge local OPN was received and forwarded to the left half Bridge, processing would proceed through step 662 to step 668 and following. This processing would result in activation

of the Enable Compare signal on line 638 to start an address comparison and causes the state machine to simultaneously activate the Delete OPN signal on line 691 in Figure 6. Activation of Delete OPN causes the Insertion/Substitution/Deletion FIFO circuit 610 to strip off the 20 bits of the OPN primitive.

The comparison in the right half Bridge comparator 402 occurs between the address of the right half Bridge local OPN previously transmitted to the left half Bridge (which was a remote OPN there and found to be of lower priority) on bus 108 and the address of the remote OPN received from the left half Bridge on bus 109. This comparison will yield the same result as it did in the left half Bridge, i.e., the right half Bridge remote OPN will be indicated as higher priority meaning the local source that generated the remote OPN received by the left half Bridge must be closed. The right half Bridge comparator 402 will activate the Preempt Local Source signal on line 670. This will cause the right half Bridge TX Port State Machine to activate the Insert CLS signal on line 684 to generate and send out a CLS on the local loop segment. This closes the local loop source node which generated the losing remote OPN at the left half Bridge.

Next, the right half Bridge must forward or unblock the remote OPN and any RRDYs, but before it can do that, it must arbitrate for and win control the local loop segment. Step 690 represents this process. First, the right half Bridge TX Port State Machine activates the Start ARB signal on line 700. This causes the FIFO circuit 610 to change any incoming CFW to ARB(0) and send them out onto the local loop segment. The local node which generated the losing OPN will have been closed so it will no longer be swallowing ARBs. When the ARB(0) arrives at each node, it is forwarded as the highest priority ARB by each node. Eventually the ARB(0) arrives at pin 1 of the right half Bridge RX port 100 and is recognized by Decode circuit 620 which then activates the ARB(0) control signal on line 702 of Figure 7 which tells the RX Port State Machine 612 that the TX port has just won arbitration. The RX Port State Machine then activates the ARB Won signal on line 704. This fact is detected by the TX Port State Machine 602 which then activates the Stop ARB signal on line 706 in Figure 6. This causes the Insertion/Substitution/Deletion FIFO 610 to stop substituting ARB(0) for incoming ARB CFWs. Then the right half Bridge TX Port State Machine activates the Insert OPN signal on line 694 which causes the FIFO circuit 610 to generate and transmit an OPN primitive. This is followed by transmission by the Insertion/Substitution/Deletion FIFO of the addresses latched in AL_PA latch 202. These addresses are always available to the FIFO circuit via bus 109 and are automatically transmitted out pin 3 in sequence following generation of an OPN.

If the remote OPN at pin 2 of the right half Bridge had been followed by any RRDYs indicating the source node on the local loop of the left half Bridge was issuing buffer credit, each RRDY would be detected by Remote

Decode circuit 604 and would cause activation of the RRDY signal on line 696. The number of time RRDY was activated is counted by the right half Bridge TX Port State Machine in an internal RRDY counter. When the state machine activates the Insert OPN signal on line 694 and sends out the addresses in sequence, it will also activate the Insert RRDY signal on line 686 a number of times equal to the number of RRDYs received and decrement the internal RRDY counter as each RRDY is sent out. This causes FIFO circuit 610 to generate and send a number of RRDY primitives after the OPN transmission equal to the number of RRDYs received from the left half Bridge. This transmission of the remote OPN and any RRDYs by the right half Bridge is symbolized by step 690 in Figure 8C.

If the winning remote OPN arriving at pin 2 of the right half Bridge TX Port was a pseudo-full-duplex OPN, the Remote Decode circuit 604 in Figure 6 of the right half Bridge detects the F7 or F8 source address and activates the Pseudo Full Duplex control signal on line 710. This causes the right half Bridge TX Port State Machine to activate the Convert Back signal on line 712 when it activates the Insert OPN signal on line 694. This causes the Insertion/Substitution/Deletion FIFO circuit 610 of the right half Bridge to delete the F7 or F8 source address and copy the destination address into the source address field to convert the pseudo-full-duplex OPN back to its original half duplex state before forwarding it onto the local loop.

Note that the losing local OPN of the right half Bridge is automatically discarded by the left half Bridge because both half Bridges automatically store OPN AL-PA addresses but delete OPN primitives when they arrive and then just regenerate the OPN if it turns out that it needs to be forwarded after winning a priority comparison. The stored addresses of a losing OPN are simply overwritten the next time another OPN arrives.

Finally, as represented by steps 699 and 701, the TX and RX ports of the left and right half Bridges go into transparent mode until a CLS is received indicating that the loop tenancy is finished. Transparent mode is explained elsewhere herein, but basically it involves passing primitive, fill words and data frames straight through the TX or RX port unchanged. Transparent mode is triggered automatically by a TX or RX port state machine when an RRDY, data frame or CLS is received which is not immediately preceded by an OPN. These events are recognized by the Remote Decode circuit 604 and the Local Decode circuit 606 in the TX port and the Decode circuit 620 in the RX ports. These decoders activate signals that tell the state machines what was received and cause the state machines to enter cause transparent mode by failing to activate any control signal to the FIFOs so no changes are made when data flows there-through and to set the state of Switches 614 and 608 so that the appropriate data path is made. After completion of step 701, processing returns to Start block 99 on Figure 8A.

Returning to the consideration of step 674 in Figure 8B, suppose the left half Bridge remote OPN won the address priority comparison. In this case path 676 is taken to step 698 in Figure 8C. Step 698 represents the process of blocking the remote OPN until a preemptive CLS can be generated and sent, arbitrating for the local loop, and, after winning arbitration, sending out the remote OPN onto the local loop along with any RRDYs received with the winning remote OPN. More specifically, this process is as follows. The remote OPN arriving at TX port pin 2 is blocked when Remote Decode circuit 604 detects the OPN and activates the OPN signal on line 632 and activates the Latch [0:1] signal on line 636. The OPN signal causes the TX Port State Machine 602 to activate the Del OPN signal on line 691 which causes the Insertion/Substitution/Deletion FIFO circuit 610 to strip off the 20 bits comprising the OPN. The Latch [0:1] signal causes the AL_PA latch 202 to store the source and destination addresses. Next, the TX Port State Machine activates the Insert CLS signal on line 684 which causes the Insertion/Substitution/Deletion FIFO circuit 610 to generate a CLS primitive and transmit it out onto the local loop segment to close the source node which generated the losing OPN. Next, the TX Port State Machine activates the Start ARB signal on line 700 which causes the Insertion/Substitution/Deletion FIFO circuit 610 to start substituting ARB(0) for any incoming CFW to gain control of the loop. When the ARB(0) CFW reaches the RX port, the Decode circuit 620 detects it and activates ARB(0) on line 702. This causes RX Port State Machine 612 to activate ARB Won on line 704 which causes TX Port State Machine to activate Stop ARB on line 706. This causes the Insertion/Substitution/Deletion FIFO circuit 610 to stop substituting ARB(0). The TX Port State Machine then activates Insert OPN on line 694 to generate an OPN and send out to the local loop via pin 3 followed by the AL-PA addresses from latch 202. If the RRDY counter in the TX Port State Machine counted any RRDYs which arrived behind the remote OPN, then the Insert RRDY signal is activated and the RRDY count is decremented and this process is repeated until the RRDY count reaches 0.

If the winning remote OPN arriving at pin 2 of the TX Port was a pseudo-full-duplex OPN, the Remote Decode circuit 604 in Figure 6 detects the F7 or F8 source address and activates the Pseudo Full Duplex control signal on line 710. This causes the TX Port State Machine to activate the Convert Back signal on line 712 when it activates the Insert OPN signal on line 694. This causes the Insertion/Substitution/Deletion FIFO circuit 610 to delete the F7 or F8 source address and copy the destination address into the source address field to convert the pseudo-full-duplex OPN back to its original half duplex state before forwarding it onto the local loop.

Note that the losing local OPN of the left half Bridge is automatically discarded by the right half Bridge because both half Bridges automatically store OPN AL-PA addresses but delete OPN primitives when they arrive

and then just regenerate the OPN if it turns out that it needs to be forwarded. The stored addresses of a losing OPN are simply overwritten the next time another OPN arrives.

Returning to the consideration of step 662 in Figure 8B, step 714 represents the process of detecting that a case 3 preemption situation has arisen. This situation is detected by the fact that the TX Arm signal on line 644 and the RX Arm signal on line 634 are both active during the same clock cycle. Each of these arming signals is activated for the clock cycle when an OPN arrives and then reset the next clock cycle. One clock cycle is all that is necessary for an address comparison or a memory access or both. When a case 3 preemption situation arises, an address comparison is performed to determine the winning OPN, and then the source node that generated the losing OPN is closed, and, if the winning OPN is a local OPN, memory is accessed to determine whether to forward the winning local OPN to the remote half Bridge or over the local bypass. The TX Port State Machine 602 does the determination of whether the case 3 situation exists by examining the states of the TX and RX Arm signals.

When a case 3 situation is found, the local and remote OPNs must be temporarily blocked, and an address comparison is started in step 716. This process starts with the TX Port State Machine activating the Enable Compare signal on line 638. This signal is also coupled to the RX Port State Machine. Activation of the Enable Compare signal on line 638 causes the TX Port State Machine to activate the Del OPN signal on line 691 in Figure 6 to cause FIFO circuit 610 to block the remote OPN from transmission until the comparison is done. In an alternative embodiment, activation of Enable Compare also causes the RX Port State Machine to activate the Del OPN signal on line 727 to cause the FIFO 616 in Figure 7 to strip the OPN primitive bits and temporarily block transmission thereof. The local OPN is automatically blocked and its address latched and converted to pseudo-full-duplex if necessary when it arrives in the preferred embodiment, as discussed elsewhere herein.

Enable Compare also causes the comparator 402 to compare the AL_PA addresses and activate one of the signals on lines 670 or 672. Step 718 represents the process of reading these comparison result signals by the RX and TX Port State Machines and branching along path 680 if the left half Bridge local OPN is higher priority, and branching along path 682 if the left half Bridge local OPN is lower priority.

Referring to Figure 8C, the processing for the situation where the left half Bridge local OPN having a lower priority will be described. The first thing that must be done is close the local source node that generated the losing OPN. This is done in step 720 where the TX Port State Machine of the left half Bridge preempts the local source node by activating the Insert CLS signal and the Start ARB signal. As a result, the Insertion/Substitution/Deletion FIFO 610 in Figure 6 generates a preemptive

CLS and transmits it out pin 3 to close the node that generated the losing local OPN. In order to forward the winning remote OPN, the left half Bridge must arbitrate for its local loop, so the Start ARB signal causes the FIFO circuit 610 to start substituting ARB(0) for incoming CFW at pin 2.

Eventually, as symbolized by step 724 in Figure 8C, these ARB(0) CFW reach the RX port and are recognized by Decode Circuit 620, as symbolized by step 722. This causes the ARB(0) signal on line 702 to be activated which causes the RX Port State Machine to activate Arb Won on line 704. The TX Port State Machine sees Arb Won and activates the Stop ARB signal on line 706 to stop the substitution of ARB(0)s, and then activates Insert OPN on line 694, and, if necessary, activates the Convert Back signal on line 712. Convert Back is only activated if the remote OPN was pseudo-full-duplex when it arrived at pin 2 thereby causing Remote Decode circuit 604 to activate the Pseudo Full Duplex signal on line 710 in addition to the OPN and Latch [0:1] signals on lines 632 and 636. Activation of Insert OPN and Convert Back causes FIFO circuit 610 in Figure 6 to generate an OPN, and convert it back to half duplex as the AL-PA addresses come by from latch 202. Alternatively, the state machine can change the F7 or F8 source address of any pseudo-full-duplex OPN back to a copy of the destination address in the latch 202 after the address comparison and before activation of the Insert OPN signal.

After forwarding the OPN and conversion back to half duplex, if the RRDY counter in the TX Port State Machine is nonzero, the TX Port State Machine activates the Insert RRDY signal on line 686 and decrements the count. This process is repeated until the RRDY count reaches zero, as symbolized by step 726. Path 728 is then taken to step 744 on Figure 8D to begin processing to enter transparent mode.

In step 744, the destination node on the local loop of the left half Bridge receives the remote OPN and responds with either an RRDY or CLS. One of these primitives propagates to pin 1 of the RX port 100 of the left half Bridge.

Step 746 represents the process of recognizing the need for transparent mode and entering it. First Decode circuit 620 in Figure 7 sees the RRDY or CLS primitive without a preceding OPN and activates either the RRDY or CLS signal on line 730 or 732, as appropriate. This causes the RX Port State Machine to control the Switch Control signal on line 622 to set Switch 614 to a 1-2 connection. Since no signal to the Insertion/Deletion/Substitution FIFO circuit 616 is activated, it acts as a transparent pipeline allowing the RRDY or CLS to propagate therethrough unchanged and pass from pin 1 to pin 2 and from there over line 54 to the remote half Bridge.

Step 748 represents the start of the process of the RX and TX ports in the right half Bridge going into transparent mode to support the loop tenancy. In step 748, the RRDY or CLS primitive transferred from the left half

Bridge arrive at the right half Bridge TX Port without a preceding OPN. The Remote Decode circuit 604 recognizes the RRDY or CLS and activates the RRDY signal on line 696 or the CLS signal on line 734. Either one of these events causes the TX Port State Machine to set Switch 608 to a 2-3 connection and control FIFO circuit 610 in Figure 6 to not alter the RRDY or CLS primitives as they pass from pin 2 to pin 3 and onto the right half Bridge local loop. The source node that originated the winning OPN sees the RRDY or CLS and responds either with a data frame or relinquishes the loop tenancy.

If a data frame is sent, the Decode circuit 620 of the RX port 100 of the right half Bridge recognizes the data frame and activates the Data Fr. signal on line 752. This causes the RX Port State Machine to set Switch 614 to a 1-2 connection because it knows that the OPN it previously sent to the left half Bridge resulted in a loop tenancy since no other OPN has been received since. The RX Port State Machine also controls FIFO circuit 616 to pass all data and primitives therethrough without change.

The data or RRDY or CLS primitives pass through the RX port of the right half Bridge and arrive at pin 2 of the left half Bridge TX port. If the arriving bits are a data frame, the Remote Decode circuit 604 activates a Data Fr. signal on line 756. If the arriving bits are an RRDY or CLS, this is recognized and the corresponding control signal is activated by the Remote Decode circuit 604. Any one of these events, not preceded by an OPN, causes the TX Port State Machine of the left half Bridge to control Switch 608 to set a 2-3 connection and control the FIFO 610 to go into transparent mode. All of the FIFO circuits 610 and 616 pass any input bits through as output bits unchanged unless one of the control signals to the FIFO is activated by the corresponding state machine. The process of the right half Bridge RX port and the left half Bridge TX port going transparent in support of the loop tenancy is symbolized by step 758 in Figure 8D. The process of going transparent by the RX and TX ports described above with reference to steps 744, 746, 748 and 758 happens whenever an OPN is forwarded and a loop tenancy is established in any of the 6 preemption rule cases described above (summarized in the Summary of the Invention section hereof) and in any case where concurrent OPNs do not occur and a normal loop tenancy is established across the two half Bridges.

A similar process of going into transparent mode occurs when an OPN establishes a purely local tenancy. In such a case after a local OPN has been forwarded to a local destination which responds with an RRDY or CLS or a data frame, the Local Decode circuit 606 sees the RRDY or CLS or data frame on pin 1 in Figure 6 coupled to the local bypass and activates one of an RRDY, CLS or Data signals on 3-conductor bus 762, activation of any one of which causes the TX Port State Machine to control Switch 608 to make a 1-3 connection and to control FIFO 610 to go into transparent mode.

The transparent state of the RX and TX ports in the left and/or right half Bridges is not altered until another OPN is received either on pin 1 of the RX port or pins 1 or 2 of the TX port.

After step 758 in Figure 8D is completed, processing returns to Start block 99 on Figure 8A.

Returning to the consideration of step 718 in Figure 8B, suppose the left half Bridge local OPN was found to be of higher priority. In this case, it is necessary to close the remote source, look up the location of the destination node of the winning local OPN, possibly convert the winning OPN back to half duplex if it is pseudo-full-duplex and its destination is local, and forward the winning OPN on the appropriate path. This process is started in step 740 on Figure 8C. This step represents the process of accessing the memory 78 in Figure 5 to determine if the destination node of the winning OPN. This is done by the RX Port State Machine by activating the Read signal on line 110 in Figure 7. When the local OPN arrived at pin 1 of the RX port, Decode circuit 620 recognized it as an OPN and activated the OPN signal on line 654 and activated the Latch [0:1] signal on line 656. The Latch [0:1] signal caused the AL_PA latch 200 to latch the source and destination addresses. If the incoming OPN was a half-duplex, Decode circuit 620 also activated the Half Duplex signal on line 658. The activation of OPN and Half Duplex causes RX Port State Machine to activate the RX Convert signal on line 660. This causes the AL_PA latch circuit 200 to substitute hex F7 or F8 for the source address in the AL_PA address stored therein. The address stored in latch 200 is presented on bus 108 to both memory 78 and comparator 402 in Figure 5. When the Read signal on line 110 is activated, the memory 78 accesses the memory location corresponding to the destination address in the AL_PA address on bus 108 and supplies the data there to the RX Port State Machine via Data bus 112. This data is a 1 if the destination is local and is a 0 if the destination is remote. This data causes RX Port State Machine to control the Switch Control signal on line 622 so as to set Switch 614 to a 1-3 connection if the destination is local, as represented by path 742 out of step 740 on Figure 8C, or a 1-2 connection if the destination is remote, as represented by path 744.

Suppose memory access step 740 concludes that the winning left half Bridge's local OPN is addressed to a local destination. In such a case, path 742 in Figure 8C leads to step 770 in Figure 8D. This step represents the process of preemptively closing the remote source node in case 3 which initiated the losing remote OPN. This is accomplished by the left half Bridge RX Port State Machine activating an Insert CLS signal on line 650 in Figure 7. This causes FIFO 616 to send a CLS primitive to the right half Bridge TX port. The right half Bridge TX port sees this CLS not preceded by an OPN and goes transparent in the manner described above in connection with the discussion of step 748. The CLS propagates to the remote source node and closes it thereby

causing the remote source node to relinquish the remote loop segment by stopping substituting its own ARB for incoming ARB fill words. This frees up the right half Bridge local loop segment to have a concurrent local loop tenancy with the left half Bridge local loop. Since the step of blocking the remote OPN at the left half Bridge accomplished in step 105 in Figure 8A automatically deleted the remote OPN, the process of discarding the remote OPN symbolized by step 770 in Figure 8D is only a failure by the TX Port State Machine of the left half Bridge to activate the Insert OPN signal on line 694 to regenerate the OPN.

Next, in step 772, the RX Port State Machine sets Switch Control signal on line 622 to cause Switch 614 to make a 1-3 local bypass connection. Although not discussed so far, the local OPN was blocked, had its AL_PA latched and automatically converted to pseudo-full-duplex if it was half-duplex when it arrived at pin 1 in Figure 7. This was done by the process of Decode circuit 620 recognizing the OPN, activating the OPN signal on line 654 and the Latch [0:1] signal on line 656 to latch the AL_PA address and activating the Half Duplex signal on line 658 if it was half duplex. This causes the RX Port State Machine to activate Del OPN on line 727 to strip the OPN bits and to activate the RX Convert signal on line 660 to automatically convert the OPN to pseudo-full-duplex if it was half-duplex when it arrived.

Because it is desirable to not send pseudo-full-duplex OPNs locally, step 776 represents the process of converting the local OPN back to half-duplex if it was half-duplex when it arrived. This is accomplished by the RX Port State Machine asserting an RX Convert Back signal on line 778. This causes the Insertion/Substitution/Deletion FIFO 616 to substitute the AL_PD destination address for the hex F7 or F8 of the pseudo-full-duplex OPN when it is regenerated. Next, the RX Port State Machine activates the Insert OPN signal on line 774 to regenerate the winning local OPN and forward it on the local segment return bypass path 60.

The winning local OPN arrives at pin 1 of the TX port from the local segment return 60 and causes the TX port to make a 1-3 connection to couple it onto the local loop segment 52 and to go into transparent mode, as symbolized by step 780 in Figure 8D. This occurs because the arrival of the OPN causes Local Decode circuit 606 to activate a Local OPN signal on line 782. Activation of Local OPN causes the TX Port State Machine to control Switch 608 to make a 1-3 connection and to not activate any signal to FIFO 610 thereby causing it to pass all input data therethrough unchanged in a transparent mode. The TX Port State Machine knows that transparent mode and a 1-3 connection is required when Local OPN is activated because this will only happen in case 3 when the local OPN has won the preemption comparison or when a local OPN is received along with no conflicting OPN and is forwarded on the local segment return 60 after a memory access. Local OPN will be activated in case 3 only after a preemption com-

parison is made, the local OPN has won, a memory access is made, the remote OPN has been discarded, and the local OPN has been forwarded on the local segment return 60.

Step 780 also represents the process of the left half Bridge going into transparent mode. This happens when the winning local OPN is forwarded onto the local loop and the destination node replies with an RRDY or CLS. The RRDY or CLS arrives at pin 1 of the left half Bridge RX port not preceded by an OPN. This causes the left half Bridge to go into transparent mode by a process similar to the process previously explained in connection with the discussion of step 758. This transparent state continues for the TX and RX ports of the left half Bridge until the loop tenancy is terminated by the source node being closed and stopping substitution of its ARB for incoming ARBs and another OPN arrives at the TX port or RX port.

Returning to the consideration of step 740 in Figure 8C, suppose that path 744 is taken. This means that a remote OPN has been received and then a local OPN was simultaneously received during the same clock cycle making a case 3 situation, and the local OPN won the address comparison and memory access indicates that the local OPN has a remote destination. In this case, step 784 is performed next to forward the winning local OPN to the right half Bridge. This is done by the left half Bridge RX port activating the Insert OPN signal on line 774. If the winning local OPN was half-duplex when it arrived and has been converted to a pseudo-full-duplex AL_PA in latch 200, no change is made as the OPN is forwarded. If the RRDY counter in RX Port State Machine is nonzero, the state machine 612 activates an Insert RRDY signal on line 786 to send out an RRDY following the OPN. The RRDY counter is then incremented, and this process is repeated to send out more RRDYs until the count reaches zero, as symbolized by step 788.

Step 790 represents the start of processing by the right half Bridge after receiving the winning OPN in case 3 from the left half Bridge. When the OPN arrives at pin 2 of the right half Bridge TX port, the Remote Decode circuit 604 activates the OPN signal on line 632 and the Latch[0:1] signal so as to store the AL_PA in latch 202. The TX port of the right half Bridge will perform the same processing as the left half Bridge when the remote OPN arrives from the left half Bridge. That process is described in Figures 8A and 8B, steps 103, 111, 105, 652, 662 and 668 and will result in the TX port of the right half Bridge doing an address comparison. Latch 200 in the RX port of the right half Bridge will still store the address of the OPN previously forwarded from the right half Bridge to the left half Bridge. The address comparison performed by the comparator 402 of the right half Bridge will compare the address in latches 200 and 202 which will yield the same result as the address comparison done by the comparator 402 of the left half Bridge. The comparator of the right half Bridge will ac-

activate the Preempt Local Source signal on line 670. This will cause the TX Port State Machine to activate the Insert CLS signal on line 684 which will send a preemptive CLS out onto the right half Bridge local loop. This closes the source node which originated the losing OPN thereby causing it to relinquish the local loop and stop substituting its own ARB for incoming ARB(0). Next, the TX Port State Machine activates the Start ARB signal on line 700 to cause FIFO circuit 610 to start arbitrating for control of the local loop by substituting ARB(0) for incoming CFW. These ARB(0)s reach the RX port, are recognized by Decode circuit 620 which activates ARB(0) which causes state machine 612 to activate Arb Won on line 704. This causes TX port state machine to activate Stop Arb on line 706 and to activate Insert OPN on line 694 to forward the OPN and send out the AL_PA address stored in latch 202. If the incoming winning OPN was pseudo-full-duplex, the Remote Decode circuit 604 will activate the Pseudo-Full-Duplex signal on line 710. This will cause the state machine 602 to activate the Convert Back signal on line 712 so that when the pseudo-full-duplex AL_PA is forwarded with the OPN, it will be converted back to half duplex.

Step 794 represents the process of activation of the Insert RRDY signal on line 686 in the right half Bridge TX port a sufficient number of times to send out on the local loop of the right half Bridge the number of RRDYs received and reduce the RRDY count to zero.

Referring to Figure 8E, step 796 represents the process which happens when the winning OPN is received by the destination node on the local loop of the right half Bridge. The destination node, when it receives an OPN will either respond with an RRDY or CLS. If the OPN is followed by RRDYs, the destination node may reply with a data frame. When the RRDY, CLS or data frame reaches the RX port, Decode circuit 620 recognizes the event and activates the appropriate signal on one of lines 752, 730 or 732. This causes the RX Port State Machine of the right half Bridge to realize that a loop tenancy has started and go into transparent mode, as symbolized by step 798 in Figure 8E. Transparent mode is entered by not activating any signal to FIFO circuit 616 and controlling switch 614 to make a 1-2 connection to connect the two since it knows from its own address comparison that the remote source node's OPN won the priority determination and the destination node is on its local loop so the two local loops have to be connected together as one large loop. Transparent mode allows the RRDY, CLS or data frame from the destination node to be forwarded from the right half Bridge RX port to the TX port of the left half Bridge. There the RRDY, CLS or data frame causes a similar set of events causing the left half Bridge TX port to go into transparent mode. When the RRDY, data frame or CLS reaches the source node, the source node replies with another RRDY, data frame or a CLS. These are received by the RX port of the left half Bridge and force it to go into transparent mode by a similar process as described above.

Receipt of a CLS by any of the TX or RX ports causes the port to go into transparent mode to forward the CLS but also signals that the loop tenancy is over so the next thing that the port can expect to receive is another OPN. To insure that when a new OPN is received, it is not automatically address compared to an old AL_PA address of an OPN from a previous loop tenancy, any TX or RX port that receives a CLS activates a Clear signal on line 800 in the case of a TX port and line 802 in the case of an RX port. Thus, whenever an OPN arrives at any RX or TX port, an address comparison can be done even if there is no conflicting OPN without fear of an incorrect result.

Returning to the consideration of step 652 in Figure 8B, suppose this test determines that TX Arm has not been recently activated. This means that a remote OPN has been received at a Bridge and no local OPN has as yet been received. This can lead to a case 4 or case 6 preemption resolution situation or a nonconflict situation where the only OPN received was the remote OPN. In such a case, path 806 is followed to step 808 in Figure 8E. Step 808 is a test to determine if the TX Arm signal on line 644 has been activated just after the remote OPN was received and arbitration has been started. If so, it means that the RX port has just received a local OPN which means that a local node has just won arbitration before the TX port ARB(0) got all the way around the local loop to the RX port. This means that the local loop is busy, so the remote OPN must be preempted. This process of closing the source on the local loop of the remote half Bridge is accomplished by step 810. The left half Bridge RX Port State Machine activates the Insert CLS signal on line 650 and insures that Switch 614 is set to a 1-2 connection. This transmits a CLS primitive to the right half Bridge TX port. This CLS primitive, not immediately preceded by an OPN, causes the TX port of the right half Bridge to go into transparent mode thereby passing the CLS to the remote source node thereby closing it and releasing the local loop of the right half Bridge for a concurrent loop tenancy, as symbolized by step 812. The CLS also causes the TX port of the right half Bridge to reset the contents of AL_PA latch 202 to an AL_PA address which will always lose in a comparison with any AL_PA of any real node in the system. This AL_PA address which will always lose in a comparison with any AL_PA of any real node in the system will hereafter be referred to as the Default AL_PA. When the TX port 102 receives the preemptive CLS not immediately preceded by an OPN, it also activates a Clr RX Latch signal on line 854 in Figure 6. Activation of this signal causes RX Port State Machine 612 in Figure 7 to clear latch 200 to the Default AL_PA by activation of the Clear signal on line 802. A Clr TX Latch signal on line 856 is activated by the RX Port State Machine any time the RX port receives a CLS to tell the TX port to clear its latch 202 to the default AL_PA by activation of the Clear signal on line 800. Any CLS not immediately preceded by an OPN will cause any RX port or TX port receiving same

to clear its AL_PA latch to this default AL_PA such that any OPN that arrives alone and is not in a conflicting OPN situation will have its AL_PA compared against the Default AL_PA and will always win.

Step 816 represents the process carried out by the RX port of beginning to handle the local OPN which either won the address priority comparison with the remote OPN or which was received alone at the RX port and was compared against the default AL_PA in latch 202. Step 816 uses the AL_PD destination address of the local OPN stored in latch 200 to access memory 78 in Figure 5 to determine the destination of the local OPN just received. This is done by supplying the latched AL_PA as the address on bus 108 to memory 78 and activating a Read signal on line 110. This causes the memory 78 to read the memory location corresponding to the address on bus 108 and return the data stored there on Data bus 112. The returned data will cause the RX Port State Machine to set Switch 614 properly in steps 825 or 826 to forward the OPN on the proper data path.

Step 825 is performed if the destination of the local OPN is remote, as determined by step 816. Step 825 represents the process performed by the RX Port State Machine 612 in Figure 7 of setting the Switch Control signal on line 622 to a state to control switch 614 to make a 1-2 connection to forward the OPN to be synthesized to the remote half Bridge. The state machine 612 then activates the Insert OPN signal on line 774 to cause the FIFO circuit 616 to synthesize and transmit the first 20 bits of an OPN out pin 2 to the right half Bridge. The address stored in latch 200 is transmitted out pin 2 after the OPN is transmitted. If the address stored in latch 200 is pseudo-full-duplex, then it is not changed in step 825 as no half duplex OPNs are allowed to be transmitted between half Bridges. If the RRDY counter in the RX Port State Machine has a nonzero count, the Insert RRDY signal on line 786 in Figure 7 is activated once for each RRDY. This causes the FIFO circuit 616 to insert an RRDY into the data stream being transmitted out pin 2 for each activation of the Insert RRDY signal. Each activation of Insert RRDY decrements the RRDY count. The Insert RRDY signal is repetitively activated until the RRDY count is zero. Processing then proceeds via path 814 to step 852 on Figure 8F.

Referring to Figure 8F, step 852 represents the process performed by the right half Bridge TX port when the OPN forwarded in step 825 arrives. In step 852, the forwarded OPN is blocked from further transmission by the FIFO circuit 610 and the AL_PA address thereof is stored in latch 202 by processes previously described. The AL_PA of the OPN just received is also compared to the AL_PA stored in latch 200 of the RX port. Since latch 200 was cleared to the default AL_PA in step 812, the AL_PA of the OPN received from the left half Bridge wins the comparison. As a result, the comparator activates the Preempt Remote Source signal on line 672 in Figure 5 to notify the RX port 100 of the right half Bridge

that the OPN received from the left half Bridge won the comparison. The RX Port State Machine ignores the activation of Preempt Remote Source and does not generate a preemptive CLS since its latch 200 has been previously cleared to the Default AL_PA which would never be the case in a preemption situation where two conflicting OPNs are still in contention. Since the only time a preemptive CLS needs to be generated is when there are still two conflicting OPNs in contention, the activation of Preempt Remote Source in this situation is ignored. The TX port 102 then starts arbitration, and if arbitration is won, forwards the OPN out onto the local loop segment after conversion back to half duplex if necessary. All this is done by processes previously described. Also, any RRDYs that followed the forwarded OPN are regenerated and transmitted out onto the local loop segment by processes previously described until the RRDY count reaches zero.

At this point, several things can happen. It is possible that the destination node of the OPN forwarded to the right half Bridge is at yet unlearned, and the OPN was forwarded because the memory 78 did not have any data indicating the destination node was on the local loop of the left half Bridge. If that is the case, the OPN will be forwarded by the left half Bridge to the right half Bridge, but the OPN will not find its destination node there. When this happens, the OPN will return to the left half Bridge. The other possibility is that the destination node will be found on the local loop of the right half Bridge. If this happens, the destination node will reply with either an RRDY, a CLS or a data frame if the OPN was followed by RRDYs issued from the source node. Which case exists is detected by step 870 in Figure 8F.

Step 870 symbolizes the process of determining whether an RRDY, CLS or a data frame has been received by the RX port or whether the same OPN forwarded onto the local loop by the TX port has been received by the RX port. Receipt of an RRDY, CLS or data frame is detected by activation by the Decode circuit 620 in Figure 7 of the RRDY, CLS or Data Fr. signals on lines 730, 732 or 752, respectively. Receipt at the RX port of the same OPN that was forwarded by the TX port is determined when Decode Circuit 620 activates the OPN signal on line 654. If this signal is activated, the OPN is known to be the same OPN as was previously forwarded because when the TX port won arbitration for the local loop, no other node on that local loop is allowed to send an OPN.

If an RRDY, CLS or data frame is received by the RX port, the process symbolized by step 872 is performed. In this process, the RX port goes transparent by processes previously described. This forwards the RRDY, CLS or data frame through the RX port to the TX port of the left half Bridge. If a CLS is received, the RX port clears its latch 200 to the Default AL_PA and activates the Clr Tx Latch signal on line 860 to cause the TX port to also clear latch 202 to the Default AL_PA.

When the RRDY, CLS or data frame reaches the

left half Bridge TX port, it goes transparent and forwards the received entity onto the local loop segment, as symbolized by step 874 in Figure 8F. If the left half Bridge TX port receives a CLS, it forwards it, and activates the Clear signal on line 800 to clear latch 202 to the Default AL_PA. The TX Port State Machine also activates the Clr Rx Latch signal on line 854 to cause the RX port to clear latch 200 to the Default AL_PA. Step 874 also represents the process of learning that the location of the destination node of the local OPN forwarded to the right half Bridge in step 825 of Figure 8E is remote, i.e., located on the local loop of the right half Bridge. The left half Bridge TX port draws this conclusion when the TX Arm signal on line 644 in Figure 7 has been activated indicating the RX port has received a local OPN and the RX Switch Pos. signal on line 664 indicates that the local OPN was forwarded to the right half Bridge. If both of these conditions are true, and the left half Bridge TX Port receives an RRDY, a CLS or a data frame on pin 2 from the right half Bridge, then the TX Port State Machine concludes that the destination node of the local OPN just forwarded is remote. Receipt of an RRDY, CLS or data frame by the TX port is detected by activation of the RRDY, CLS or Data Fr. signals on line 696, 734 or 756, respectively. When one of these signals is activated in the circumstances defined above, the TX Port State Machine activates the Clr Mem signal on line 960 in Figure 6. Activation of Clr Mem causes the RX Port State Machine to activate the Clear signal on line 930. This clears the memory location associated with the destination address portion of the address on bus 108. Since a Select TX signal on line 950 is not activated, an internal multiplexer inside memory 78 remains in the default condition connecting the address on bus 950 to the address port of the memory array in memory 78 of Figure 5. This causes a logic 0 to be written into the forwarding table location associated with the local OPN's destination node indicating it is remote relative to the left half Bridge.

As symbolized by step 876 in Figure 8F, the entity passed through the TX port of the left half Bridge will cause the source node to respond with another RRDY, a data frame or a CLS. This RRDY, data frame or CLS arrives at the left half Bridge RX port and causes it to go transparent and pass the entity to the TX port of the right half Bridge. The TX port of the right half Bridge then goes transparent and passes the received entity through onto the local loop segment of the right half Bridge. Of course, if the source node receives a CLS, it stops substituting its arb primitive for incoming arb primitives and relinquishes the loop tenancy. If the source node responds with an RRDY, CLS or data frame, the RX port of the left half Bridge forwards them to the TX port of the right half Bridge which goes transparent by processes previously explained and forwards the entity received to the destination node. The loop tenancy continues in this manner until either the source node or destination node transmits a CLS causing the source node to relinquish the

loop tenancy. Processing then proceeds via path 877 to step 879 on Figure 8G which vectors processing back to Start on Figure 8A.

Returning to the consideration of step 870 in Figure 8F, if the RX port of the right half Bridge receives an OPN after the local OPN was forwarded from the left half Bridge to the right half Bridge and forwarded onto the right half Bridge local loop, a case 5 situation has arisen. In a case 5 situation, a local OPN was received, forwarded to the other half Bridge and is returned to the half Bridge which forwarded it. This means that the destination node is not on the local loop of the half Bridge to which the OPN was forwarded or does not exist. In this situation, step 870 vectors processing to step 880. Step 880 symbolizes the process of the right half Bridge going transparent and forwarding the OPN back to the left half Bridge by processes previously described. The RX port determines that the OPN it just received is the same OPN that the right half Bridge TX port previously received by checking the status of the RX Arm signal on line 634 in Figure 7 and the status of the Equal signal on line 850 in Figure 7 when the OPN signal on line 654 goes active. When the OPN was received by the TX port, its AL_PA was latched into latch 202. When the OPN is received by the RX port, its AL_PA is latched in latch 200, and the TX Arm signal on line 644 in Figure 7 is activated. This causes the TX Port State Machine to activate the Enable Compare signal on line 638. This causes the comparator 402 in Figure 5 to compare the addresses in latches 202 and 200 and activate the Equal signal on line 850. When the RX port finds that the RX Arm and Equal signals are both active when OPN on line 654 goes active, the RX port knows that the OPN just received has the same AL_PA as the OPN previously received by the TX port and must be forwarded back to the left half Bridge. The RX Port State Machine in this case implements the required transparency by not activating the Del OPN signal on line 727 in Figure 7 in response to activation of the OPN signal on line 654 in Figure 7. This allows the OPN to pass through the FIFO pipeline and propagate back to the left half Bridge. The same is true for any following RRDYs.

Step 882 represents the process which happens when the OPN forwarded by the right half Bridge RX port arrives as a remote OPN at the left half Bridge TX port. In step 882, the left half Bridge TX port detects the arrival of the remote OPN, stores its AL_PA in latch 202 and activates the Del OPN signal on line 691 in Figure 6 to block the transmission of the OPN until the possible preemption situation is resolved. The possible preemption situation is signalled by an active TX_Arm signal on line 644. This signal was activated by the RX port when it received the local OPN which was forwarded to the right half Bridge and which later returned. To resolve the possible preemption situation, the TX Port State Machine activates the Enable Compare signal on line 638 in Figure 6. This causes comparison of the AL_PA of the OPN just received and stored in latch 202 against

the AL_PA stored in latch 200. The AL_PA stored in latch 200 represents the AL_PA of the OPN previously forwarded to the right half Bridge from the left half Bridge RX port. This comparison results in activation of the Equal signal on line 640 in Figure 6. This causes the TX Port State Machine to conclude the remote OPN just received is the same local OPN which was previously forwarded to the right half Bridge. As a result, the TX Port State Machine activates the Start Arb signal on line 700. This causes the Insertion/Substitution/Deletion FIFO circuit 610 to start substituting ARB(0)s for incoming fill words. Since these are the highest priority arb primitives, they pass through all the nodes on the local loop of the left half Bridge and arrive at the left half Bridge RX port. When the ARB(0) arrives at the RX port, the Decode circuit 620 activates the ARB(0) signal on line 702 in Figure 7. This causes the RX Port State Machine 612 to activate the Arb Won signal on line 704. This causes the TX Port State Machine 602 to activate the Insert OPN signal on line 694 thereby causing an OPN to be generated and to propagate onto the local loop segment of the left half Bridge. Then the TX Port State Machine activates the Insert RRDY signal on line 686 enough times to cause the number of trailing RRDYs received to be forwarded and the RRDY count reduced to zero. The TX port also activates the RX Arm signal when the remote OPN is received to tell the RX port that it has received an OPN. Processing then proceeds via path 884 to step 886 on Figure 8G.

There is a possibility that the destination node of the OPN just forwarded is not on the local loop of the left half Bridge either. In this case the OPN forwarded by the left half Bridge TX port will propagate through the local loop of the left half Bridge and arrive at the left half Bridge RX port. Step 888 determines if an OPN has been received by the RX port. If so, the RX Port State Machine activates the Auto_Preempt signal on line 646. This causes the TX Port State Machine to activate the Insert CLS signal on line 684. This sends a CLS out onto the local loop of the left half Bridge to close the source node to relinquish the loop tenancy. This process is symbolized by step 900 in Figure 8G. Processing then returns to Start on Figure 8A.

Returning to the consideration of test 886 in Figure 8G, assuming the destination node is located on the local loop of the left half Bridge, the OPN forwarded by the TX port arrives at this destination node. As symbolized by step 888, the destination node receives the OPN and any following RRDYs and responds by transmitting an RRDY, a CLS or a data frame. Step 890 represents the process performed by the left half Bridge RX port when this RRDY, CLS or data frame arrives there. The RX port recognizes the arrival of an RRDY, CLS or data frame by activation of one of the RRDY, CLS or Data Fr. signals on lines 730, 732 or 752, respectively, in Figure 7. When the RX Port State Machine 612 sees activation of one of these signals coupled with an active Equal signal on line 640 and an active RX Arm signal on line 634,

it concludes that the destination node of the OPN it previously forwarded to the right half Bridge must be on the local loop of the left half Bridge. Therefore, in step 890, the RX Port State Machine activates the Set signal on line 116 which causes the memory location in memory 78 in Figure 5 corresponding to the destination node address of the OPN just received by the TX port to be written with data indicating this destination node is on the local loop of the left half Bridge. The RX port then set switch 614 to a 1-3 connection to bypass the right half Bridge, and goes into transparent mode to pass whatever entity was received from the destination node therethrough. By setting the 1-3 connection, the RX port passes the entity received from the destination node to the TX port via the local segment return 60. Also, if a CLS was received, the RX port activates the Clear signal on line 802 and the Clr Tx Latch signal on line 860 to cause both latches 202 and 200 to be cleared. Processing then proceeds back to Start on Figure 8A.

Returning to the consideration of Figure 8E, if step 816 finds that the destination of the local OPN is local, step 826 on Figure 8E is performed. This step represents the process wherein the RX Port State Machine controls Switch 614 to make a 1-3 connection and to activate the Insert OPN signal on line 774. This causes an OPN to be generated and sent out on the local segment return 60. Whatever address is stored in latch 200 is sent out behind the OPN. The RX Port State Machine also activates a Bypass signal on line 114 to tell the TX port that the OPN that it is sending on the Local Segment Return 60 is destined for a local node and the TX port should go transparent. If the address stored in the latch 200 is pseudo-full-duplex, the RX Port State Machine activates an RX Convert Back signal on line 778. This causes the Insertion/Substitution/Deletion FIFO circuit 616 to alter the AL_PA read out of the latch 200 as it streams by and convert the hex F7 or hex F8 source address to a copy of the AL_PD destination address so as to convert the OPN back to its original half-duplex form. If the RRDY counter in the RX Port State Machine has a nonzero count stored therein, the state machine 612 activates the Insert RRDY signal on line 786 so as to cause the FIFO circuit 616 to generate an RRDY and send it out behind the OPN. This process is repeated until the RRDY count reaches zero.

After the local OPN has been forwarded, processing proceeds via path 813 to step 830 on Figure 8F. In this step, the TX port receives the OPN on the local segment return. The TX Port State Machine learns of the arrival of this OPN when a Local Decode circuit 606 detects the OPN and activates a Local OPN signal on line 782. This causes the TX Port State Machine to check the status of the Bypass signal on line 114. If Bypass is active, the TX port goes into transparent mode, controls Switch 608 to make a 1-3 connection, and passes the OPN onto the local loop without arbitrating or blocking the transmission therethrough. In an alternative embodiment, no Bypass signal is used, and whenever the Lo-

cal OPN signal on line 782 is activated, the TX Port State Machine assumes that the RX port has verified the destination is local and automatically goes into transparent mode and makes a 1-3 connection.

At this point, the OPN may reach the destination node and cause it to issue an RRDY, a data frame or a CLS. In addition, the destination node may have been unplugged or failed thereby causing the OPN to return to pin 1 of the RX port. Further, the destination node may respond with an RRDY which travels along the local loop segment and then reaches the source node which originated the OPN and which is closer to the RX port than the destination node. This source node may itself respond with a CLS. Test 832 determines which is the case. If the RX port receives an RRDY or CLS or a data frame, processing vectors to step 834. Receipt of an RRDY or CLS or Data is detected by activation of the RRDY signal on line 730 or activation of the CLS signal on line 732 or activation of the Data FR signal on line 752 in Figure 7 by Decode circuit 620. If any of these signals is activated, step 834 symbolizes the process wherein the RX port goes transparent on the 1-3 connection and passes the primitive or data through to the local segment return 60. In addition, if a CLS is received, the RX Port State Machine clears latch 200 to the Default AL_PA by activating the Clear signal on line 802. Either the RRDY or CLS or data frame is passed to the TX port. This is detected by the Local Decode circuit 606 which activates either the RRDY, CLS or Data signal, as appropriate, on 3-conductor bus 762. Since the TX port went into transparent mode in step 820, the primitive or data frame passes straight through to the destination node. If the CLS signal on bus 762 is activated, the TX Port State Machine activates the Clear signal on line 800 to clear latch 202 to the Default AL_PA address, as symbolized by step 836. Finally, processing returns to Start on Figure 8A.

If step 832 determines that the RX port received an OPN, that OPN could only be the same OPN that was previously transmitted over the local segment return by the left half Bridge RX port to the left half Bridge TX port. This is because the remote OPN from the right half Bridge was previously preempted and no other node could win arbitration and gain permission to send until the source node that originated the local OPN relinquishes the loop. Receipt of the same OPN previously forwarded on the local bypass 60 by the RX port also means that the destination node was either removed from the left half Bridge local loop or has failed or the memory contains incorrect data and the destination node is actually remote. This is true because the RX port would not have forwarded the local OPN onto the local segment return 60 unless the data in memory 78 indicated that the destination was local. RX port 100 will know that it received the same OPN back by virtue of activation of the OPN signal on line 654 in Figure 7 immediately after forwarding the OPN onto local segment return 60. In alternative embodiments, the Local Decode

circuit 606 could activate a signal similar to the Latch [0:1] signal and the AL_PA latch 202 would also be connected to the local return segment 60 so as to latch the AL_PA address of the OPN received on the local segment return 60. In such a case, the receipt of the OPN by the RX port would cause the comparator to compare the address in the latch 200 in Figure 7 to the address in latch 202 in Figure 6. This comparison would cause activation of an Equal signal on line 850. When the same local OPN is received back at the RX port, step 838 in Figure 8F is performed to update the memory to reflect the fact that the destination node is now no longer local and is either remote or unknown. The left half Bridge RX port then activates the Insert CLS signal on line 650 so as to send a CLS to the local source node and close it thereby extinguishing the loop tenancy. Processing then returns to Start on Figure 8A.

Returning to the consideration of step 808 in Figure 8E, if TX Arm did not go active just after the remote OPN was detected, then path 910 is taken to step 912. In this step, the TX port starts substituting ARB(0) for incoming arb fill words so as to obtain control of the local loop. When the ARB(0) primitive reaches the RX port, arbitration has been won, and the RX port activates the Arb Won signal on line 704. This causes the TX Port State Machine to activate the Insert OPN signal on line 694 so as to regenerate the OPN and to place the AL_PA of the remote OPN into FIFO 610 for transmission onto the local loop segment 52 behind the OPN. Also, if the RRDY count is nonzero, the TX port state machine activates the Insert RRDY signal on line 686 in Figure 6 a sufficient number of times to insert the number of RRDYs received behind the OPN and reduce the RRDY count to zero.

The OPN then propagates through the local loop and may or may not find its destination node there. If the destination node is found, it will reply with an RRDY, CLS or data frame which will propagate to the RX port. If the destination node is not found, the OPN will propagate to the RX port. Test 914 determines whether the RX port has received an OPN, an RRDY, a CLS or a data frame.

If test 914 detects an OPN, a case 6 situation has arisen where a remote OPN arrives, is forwarded onto the local loop and returns to the RX port. In this case, the arbitration process started by the TX port must be stopped, the OPN must be forwarded to the right half Bridge, and the left half Bridge TX and RX ports must go into transparent mode and the memory must be updated to indicate the destination node is either remote or unknown relative to the left half Bridge. Accordingly, processing is vectored to step 916 where the memory 78 is updated to indicate the destination node is either remote or unknown. This is done by the RX Port State Machine when the OPN signal on line 654 in Figure 7 and the Equal signal on line 640 and the RX Arm signal on line 634 are all active. This results in activation of the Clear signal on line 930. This causes the memory location in memory 78 corresponding to the address of the

OPN in latch 200 to be cleared indicating the destination address is either remote or unknown. In some embodiments, the OPN received by the RX port will be blocked in step 916 by processes previously described and then forwarded to the right half Bridge later when all of the above identified signals have been activated.

Processing is then vectored to step 918 where the OPN is forwarded to the right half Bridge and the left half Bridge transitions to transparent mode. In the preferred embodiment, the RX Port State Machine detects activation of the OPN signal on line 654 in Figure 7 and the Equal signal on line 640 and the RX Arm signal on line 634 all in the same clock cycle and automatically transitions into transparent mode. When these signals are all active, it means that the same OPN forwarded from the left half Bridge TX port has arrived at the RX port not having found its destination node. This means that the OPN must be forwarded to the right half Bridge. Thus, step 918 represents the process wherein the RX Port State Machine set Switch Control signal on line 622 to set Switch 614 to make a 1-2 connection, and fails to activate the Del OPN signal on line 727. This causes the OPN to pass through FIFO circuit 616 unhindered and be transmitted out pin 2 to the right half Bridge. In this situation, step 918 also represents the process carried out by the RX Port State Machine of activating a Go Transparent signal on line 922. This is detected by the TX Port State Machine which responds by not activating any signal coupled to the Insertion/Substitution/Deletion FIFO circuit 610. This causes any primitive or data frame entering the TX port at pin 2 to pass through the FIFO unhindered. Step 918 also represents the process carried out by the Insertion/Substitution/Deletion FIFO circuit 610 of continuing to substitute ARB(0)s for incoming fill words thereby preventing nodes on the left half Bridge local loop from winning arbitration. This reserves the local loop of the left half Bridge for use in communication of primitives between the source node and destination node even if both nodes are on the local loop of the right half Bridge. Note that the Case 6 situation typically arises in only one situation: where both the source node and destination node are on the right half Bridge local loop (or the left half Bridge local loop for the mirror image situation) but the source node is closer to the RX port than the destination node and the location of the destination node is as yet unlearned.

Processing of the forwarded OPN by the right half Bridge in Case 6 is started by performing steps 103, 111 and 105 in Figure 8A and steps 652, 662, 668 and 674 in Figure 8B. In test 674, the comparison by comparator 402 in Figure 5 in the right half Bridge results in activation of the Equal signal on line 640. This results in vectoring of processing via path 936 to step 938 on Figure 8G.

Step 938 represents the process performed by the TX port of the right half Bridge in blocking the remote OPN and any following RRDYs by processes previously described. Step 938 also represents the process per-

formed by the right half Bridge 1 detecting that the remote OPN just received is the same OPN as the local OPN previously forwarded to the left half Bridge by the RX port. This recognition happens when OPN on line 654 in Figure 6 goes active while the TX Arm signal on line 644 is active followed by activation of the Equal signal on line 640. The TX port automatically activates the Enable Compare signal on line 638 in Figure 6 when the remote OPN arrives. This causes comparison between the AL_PA addresses in the latches 200 and 202, and, since the RX port latch contains the AL_PA of the local OPN previously forwarded to the left half Bridge, the Equal signal is activated by the comparator. Activation of Equal causes the TX Port State Machine to activate the Start Arb signal on line 700. This causes the FIFO circuit 610 to start substituting ARB(0) for incoming fill words. When an ARB(0) reaches the RX port of the right half Bridge, the RX Port State Machine sees activation of ARB(0) on line 702 in Figure 7 which causes activation of the Arb Won signal on line 704. This activation of Arb Won causes the TX Port State Machine to activate the Stop Arb signal on line 706 to stop the arbitration process followed by activation of the Insert OPN and Insert RRDY signals on lines 694 and 686, respectively to transmit the OPN and any following RRDYs out onto the local loop segment.

Next, step 940 on Figure 8G is performed which represents the process wherein the destination node receives the forwarded OPN and any following RRDYs and replies with an RRDY, a data frame or a CLS.

Step 942 represents the process of the source node receiving the RRDY, CLS or data frame and reacting appropriately. The source node may send a data frame if it receives an RRDY, or may transmit another RRDY if it receives a data frame and has sufficient buffer capacity to receive another frame. Or, if the source node receives a CLS, it will stop substituting its arb fill word for incoming fill words thereby relinquishing the loop. The RX port watches the fill words during a loop tenancy and clears its AL_PA latch to the Default AL_PA when the fill word changes indicating that the loop tenancy is finished thereby preparing itself for receipt of the next OPN and comparison of the AL_PA of the OPN against either the Default AL_PA or the AL_PA of a conflicting OPN.

If the source node responds with either an RRDY or data frame, the RX port detects this fact in step 944 when either the RRDY signal on line 730 in Figure 7 or the Data Fr. signal on line 752 is activated by the Decode circuit 620. When either one of these signals is activated with the RX Arm signal active on line 634 and the Arb Won signal having been previously activated, the RX Port State Machine concludes that the destination and source nodes of the current loop tenancy are both on the right half Bridge local loop. The RX Port State Machine then writes data to memory 78 to memorize this fact by activating the Set signal on line 116 thereby causing the memory location in memory 78 corresponding to the AL_PA of the OPN stored in latch 200 to be written

with a logic 1. The RX Port State Machine then sets Switch 614 to a 1-3 bypass connection to keep the loop tenancy local and goes into transparent mode so that incoming primitives, data frame and fill words pass therethrough on the 1-3 path unchanged and undelayed except for the FIFO delay. The RX Port State Machine then also activates the Go Transparent signal on line 922 and the

Bypass signal on line 114 to cause the TX port to set Switch 608 to a 1-3 connection to connect the local segment return 60 to the local loop segment 52 and also enter the transparent mode, as symbolized by step 946. Processing then returns to Start on Figure 8A.

After step 918 is performed, processing returns to Start on Figure 8A.

Referring again to test 914 in Figure 8E, if the RX port receives an RRDY or CLS or data frame, processing vectors to step 948. Step 914 is reached when a remote OPN has been received, no previous local OPN has been received and TX arm does not go active just after the remote OPN has been received and before the TX port won arbitration in response to receiving the local OPN. In this situation, the remote OPN is forwarded out onto the local loop segment in step 912, and when the RRDY or CLS or data frame is received, the RX port knows that the destination node of the remote OPN is on the local loop. Accordingly, step 948 represents the process of learning that the location of the destination node of the remote OPN is on the local loop by writing data into the memory location of memory 78 in Figure 5 that corresponds to the AL_PA of the remote OPN, said data indicating that the destination node is local. The way this is done is by activating a Select TX signal on line 950 whenever the RX Arm signal is active and the Decode circuit 620 in Figure 7 activates the RRDY or CLS or Data Fr. signals on lines 730 or 732 or 752, respectively. The TX Arm signal is active only for the clock cycle in which the remote OPN is received since the conflicting OPN preemption comparisons and generation of preemptive CLSs and synthesis of the winning OPN for forwarding takes only one clock cycle. The memory 78 has an address multiplexer internal thereto which has two inputs and one output coupled to the address port of the memory array (not shown). One input is the default input and is coupled to bus 108 to supply the AL_PA address from the RX port latch 200 to the multiplexer output whenever the Select TX signal is not active. The other input is coupled to bus 109. When Select TX is active, the address from the TX port latch 202 on bus 109 is coupled to the multiplexer output and the memory array address port. Step 948 represents the process of activating the Select TX signal and activating the Set signal on line 116. This causes the memory location corresponding to the destination address portion of the AL_PA address of the remote OPN on bus 109 to be set to logic 1 indicating that this destination address is local.

Processing next proceeds to step 952 which represents the process of the RX port going into transparent

mode and forwarding the RRDY or CLS to the right half Bridge. This is implemented by the RX Port State Machine not activating any control signal to the FIFO circuit 616 and controlling the Switch Control signal on line 622 to set Switch 614 to a 1-2 connection. These two actions cause any data entering the RX port at pin 1 to pass through the FIFO unchanged and be transmitted out pin 2. The forwarded RRDY or CLS arrive at the TX port of the right half Bridge unpreceded by an OPN and cause it to go into transparent mode thereby forwarding these primitives to the source node which responds with another RRDY, data frame or CLS. The data or primitive transmitted by the source node causes the RX port to go transparent and pass the data or primitive to the TX port of the left half Bridge which goes transparent and passes them to the destination node. This process continues until a CLS is transmitted which causes the loop to be relinquished and the latch of any RX or TX port is passed through to be cleared to the Default AL_PA. Processing then returns to Start on Figure 8A. in step 954.

Referring to Figure 10, there is shown a flow chart of the portion of the software that controls processing by the Transaction Server 16 in Figure 1 to perform mirrored write transactions. The specific location in the software architecture of the Transaction Server for this software is not critical to the invention. For example, the software of Figure 10 can be part of an application program at an upper layer or can be part of the control software managing an FC-AL interface circuit which couples the processing circuitry of the Transaction Server to the FC-AL medium 10. Refer not to Figures 10 and 1. In Figure 10, test 1000 causes the Transaction Server 16 in Figure 1 to determine if a write request has been made to the primary bank of disk drives 12 and 14. If not, processing returns to the main loop of whatever program incorporates the process of Figure 10 to continue processing, as symbolized by step 1002. Test 1000 may be periodically executed as a subroutine or may be executed in response to a hardware or software interrupt service routine which is executed when a write transaction is performed.

If a write transaction has been detected in test 1000, the processing of block 1004 is carried out. This process generates an OPN addressed to the appropriate hard disk in the primary bank of disk drives and writes the data of the requested write transaction thereto in the ensuing loop tenancy started by the OPN. Next, the process of block 1006 is performed wherein another OPN is generated and addressed to the mirrored hard disk drive of the Secondary Disk Drive banks 32 and 34 (sometimes referred to herein as the Mirrored Storage Disk Drives) on the Secondary Loop 26. The purpose of this OPN is to attempt to start a loop tenancy between the Transaction Server 16 and the mirrored disk drives 32 and 34 across Bridge 28. When the OPN addressed to the Secondary Disk Drives 32/34 or the Backup and HSM Server 30 reaches the Bridge 28, it will be forwarded onto the Secondary Loop 26 if the Secondary Loop

is not already busy with a loop tenancy involving backing up data from the Secondary Disk Drives to the Archival Storage Devices 38 or 40. If the OPN from the Transaction Server 16 arrives simultaneously at the Bridge 28 with an OPN on the Secondary Loop 26 such as an OPN from Backup and HSM server 30 to the Mirrored Storage Disk Drives, the Bridge 28 is programmed to kill the OPN on the Secondary Loop 26 by generating a preemptive CLS and sending it around the Secondary Loop 26. Then, the OPN from the Primary Loop 10 is forwarded onto the Secondary Loop to start the loop tenancy for the mirrored write transaction. In this case, the Bridge 28 couples the Primary Loop 10 to the Secondary Loop 26 for the duration of the mirrored write loop tenancy to make a single FC-AL. The Backup and HSM Server 30 can be programmed to retry the OPN for the archive or other HSM transaction later.

If the Secondary Loop is already busy with a loop tenancy when the OPN directed to the Mirrored Storage Disk Drives 32/34 arrives, the Bridge 28 will generate a preemptive CLS and send it back around Primary Loop 10 to Transaction Server 16. This fact is detected by test 1008 in Figure 10. Test 1008 determines if a CLS came back in response to the OPN transmitted to the Secondary Loop by the Transaction Server before the mirrored write transaction data was successfully recorded on the Mirrored Storage Disk Drives 32/34. A CLS may be returned to the Transaction Server 16 either if the Secondary Loop is busy when the OPN from the Transaction Server arrives at Bridge 28, or if the OPN from the Transaction Server is forwarded onto the Secondary Loop, but the destination node has no buffer capacity to receive data and issues a CLS. In either case, if the answer to the question posed by test 1008 is yes, the process of block 1010 is performed. This block represents the process of buffering the data and trying the mirrored write transaction again later. This process can be carried out either in the transaction server 16 or the Bridge 28. If carried out in the transaction server, the data to be written in the mirrored write transaction is stored in memory somewhere and the transaction server periodically or occasionally generates an OPN addressed to the Mirrored Storage Disk Drives and transmits it on primary loop 10 to the Bridge 28 until the OPN is forwarded, the data is completely written onto the Mirrored Storage Disk Drives and the loop tenancy is terminated. If the process is carried out in the Bridge, when the Bridge sends a preemptive CLS and the CLS is detected in step 1008, the transaction server sends the data to be written in the mirrored write transaction to the Bridge with an instruction to store the data and write it to the Mirrored Storage Disk Drives as soon as the loop tenancy on the Secondary Loop 26 is finished. In this embodiment, the Bridge acts as a node and has sufficient storage capacity to store the mirrored write data and has sufficient intelligence to receive, decode and act upon the command from the transaction server to carry out the mirrored write transaction as soon as the Secondary Loop

is free.

If test 1008 determines that a preemptive CLS did not come back immediately in response to the OPN generated by the Transaction Server directed to a node on the Secondary Loop, then the process of block 1012 is performed. This block represents the process of waiting for an RRDY to come back from the secondary loop across the Bridge to the Transaction Server 16. When the RRDY does arrive, the Transaction Server knows that the Bridge has coupled the Primary Loop to the Secondary Loop and it proceeds to carry out the mirrored write transaction via a standard loop tenancy between itself and the Mirrored Storage Disk Drives. Thereafter, processing returns to the main loop, as symbolized by block 1014.

Referring to Figure 11, there is shown a block diagram of an alternative embodiment which employs the teachings of the invention but which does not use a separate Bridge. The embodiment is similar to the embodiment of Figure 1, and all elements having reference numbers which are the same as reference numbers in Figure 1 are the same elements as the element in Figure 1 with the corresponding number and serves the same purpose in the combination. Instead of using a separate Bridge, the Bridge function is carried out by software in the Transaction Server 1200 and the HSM server 1202 and by a third Fibre Channel Arbitrated Loop or other local area network 1204. In the embodiment of Figure 11, each of the servers 1200 and 1202 have two Fibre Channel Arbitrated Loop interface cards (hereafter FC-AL Interface). In the case of Transaction Server 1200, one of these FC-AL Interfaces couples server 1200 to the Primary Loop 10, while the other FC-AL Interface couples the Transaction Server 1200 to the third FC-AL 1204.

Likewise, HSM server 1202 has two FC-AL Interfaces: one couples HSM server 1202 to the Secondary Loop 26; and the other couples the HSM server 1202 to the third loop 1204.

The Transaction Server 1200 includes software that performs the bridging function as follows. During normal read transactions, the bridging software in the Transaction Server sends an OPN to the disk drives 12, 14 and carries out a normal SCSI or IP read transaction in which neither the third loop 1204 nor the HSM server 1202 need be involved. On a write transaction, the Bridge software in Transaction Server 1200 does a conventional OPN addressed to disk drives 12, 14 and then does a conventional SCSI or IP write transaction to the disk drive so opened. Next, after completion of this write, the bridging software generates an OPN addressed to the HSM server 1202 or mirrored disk drives 32, 34. This OPN is sent by the bridging software to the FC-AL Interface card which couples Transaction Server 1200 to the third loop 1204 which transmits it on the third loop. The bridging software in the HSM server 1202 is constantly monitoring the third loop for traffic. When it sees an OPN addressed to the HSM server 1202 or any other

node on the Secondary Loop 26, the HSM Server bridging software transfers the traffic to the appropriate node by copying the traffic onto the Secondary Loop 26 when necessary. This is done by copying the traffic from the third loop 1204 using the first FC-AL Interface and sending it to the second FC-AL Interface for transmission to the destination node on the secondary loop 26 (unless the traffic is directed to the HSM server itself, in which case the bridging software does an interprocess transfer of the traffic to whatever other HSM software process needs to receive the traffic).

Likewise, the bridging software in the Transaction Server 200 is constantly monitoring the third loop for traffic addressed to a node on the Primary Loop 10. When such traffic is found, it is copied from the second FC-AL Interface circuit coupled to the third loop 1204 and transferred to the second FC-AL Interface coupled to the Primary Loop 10 from which it is transmitted over the Primary Loop to its destination node. These inter-Interface circuit transfer paths are maintained by the bridging software until a CLS primitive has been detected by the bridging software which indicates the connection is no longer needed. Generally, the only traffic travelling from the Secondary Loop 26 to the Primary Loop 10 will be messages from the HSM server 1202 to update its file system by marking certain blocks, sectors and tracks as archived and available for overwriting, to flush its cache storage of file system data in some circumstances to maintain file system accurateness, and archival data read by the Transaction Server from the archival storage devices.

The HSM server 1202 does its HSM functions using only the Secondary Loop 26, except for message traffic that is sent to the Transaction Server 1200 to cause it to update its file system properly to reflect the fact that certain data has been archived, etc. By performing the HSM functions on the Secondary Loop 26 only and not involving the Primary Loop, the principle advantage of the invention can be achieved, i.e., online backups of the entire mirrored storage can be achieved simultaneously with continued read processing by the Transaction Server 1200 without slowing down the read transaction processing on the Primary Loop.

Referring to Figure 12, there is shown another alternative embodiment of the system of Figure 1 using a single transaction processor 1210. This single transaction processor is coupled to the Primary Loop 10 through a first FC-AL Interface card 1212 and is coupled to the Secondary Loop 26 through a second FC-AL Interface 1214. This single transaction processor 1210 performs both the transaction processing and the HSM functions using dual, parallel processing processors 1216 and 1218. Use of two separate processors prevents the normal transaction processing from being slowed down by processing required by HSM functions. The transaction processor 1210 is coupled to a plurality of workstations 22 and 24 via an Ethernet or other local area network arrangement 18/20/21/23 and a suitable interface card

1220 that can couple the local area network 18 to the internal bus 1222 of the transaction processor 1210. Likewise, the transaction processor is coupled to the backup storage devices 38 and 40 via a SCSI bus 36 and a suitable SCSI bus interface circuit 1221 to couple the SCSI bus 36 to the internal bus 1222. Typically, the internal bus 1222 is a PCI bus.

Bridging software in execution on the two processors 1216 and 1218 operates as described above to Bridge any necessary traffic between the primary and Secondary Loops 10 and 26. This cross traffic between the two loops will be traffic necessary to carry out mirrored write transactions to mirrored storage disk drives 32, 34, bring over archived data from the backup storage devices 38 and 40 and messages informing the transaction processor 1216 to update its file system data structure when the HSM processor 1218 archives old and/or infrequently used data from the mirror disk drives 32, 34 to the backup/archive storage devices 38 and 40. Other than the aforementioned cross traffic between the two loops, the primary and second loops operate independently and simultaneously with regard to read transactions on the Primary Loop and HSM backups on the Secondary Loop.

Although, no internal volatile memory is shown in transaction processor 1210, those skilled in the art understand that such memory exists in the system and it too shares the internal PCI bus 1222 if it is internal to the processor 1210. Thus, since all internal traffic to and from memory, and all transaction processing traffic between processor 1216 and first FC-AL Interface 1212, and all HSM traffic between processor 1218 and second FC-AL Interface 1214 and all traffic to and from the interfaces 1220 and 1221 must share the internal bus 1222, the internal bus 1222 can be a bottleneck. As such this embodiment is the least preferable of all the embodiments disclosed herein.

Although the invention has been disclosed in terms of the preferred and various alternative embodiments disclosed herein, all apparatus and processes which isolate the HSM processing on a separate local area network from the main transaction processing, uses one or two separate processor for carrying out the transactions on each network and which uses a Bridge or bridging process to isolate the two networks and only selectively cross couple the two networks for necessary cross traffic will fall within the teachings of the invention. For example, the primary and Secondary Loops in the embodiments of Figures 1, 6 and 7 could be replaced by ATM networks or any other local area network topology and protocol, and the higher level SCSI read and write protocols could be replaced with any other protocol which the disk drives and archival backup storage devices understand such as IPI or prior art proprietary protocol that existed before the SCSI industry standard emerged. A key concept that defines, among other things, the genus of the invention is in using a Bridge or bridging type process to cross connect the two networks

only selectively where necessary to carry out the mirrored write transactions and keep the file systems of the two transaction processors consistent but otherwise keeping the networks isolated from each other so that HSM transactions on one network do not tie up transaction processing on the other network.

In other alternative embodiments, the system of Figure 1 may be modified by removing the HSM server 30 and putting HSM server and backup functionality on the Transaction Server. In such an embodiment, the secondary backup devices 38 and 40 will connect directly to the secondary FC-AL 26. In addition, the mirrored storage hard disk drives 32 and 34 have intelligence sufficient to receive Copy commands from the Transaction Servers designating the blocks to be copied, and to initiate a loop tenancy by transmitting an OPN to the target secondary backup device of devices 38 and 40 and transmitting the designated blocks to the secondary backup device. Also, the backup and HSM server 30 can be kept informed by the Transaction Server 16 by management and control messages across Bridge 28 so that the HSM server can double as a standby Transaction Server in case of failure of the Transaction Server 16. Such management and control messages includes messages indicating file system updates indicating the time and date each file is updated, which files are open etc. and any other messages necessary for the HSM server to step in and take over for the Transaction Server 16 if it fails.

Claims

1. An apparatus comprising:

a primary memory;
a mirrored memory;
a backup/archival storage device;
means for carrying out read and write transactions with said primary memory over a first local area network and, for every write transaction to said primary memory, carrying out a write transaction of the same data to said mirrored memory via a second local area network, and for carrying out read and write transactions between said mirrored memory and said backup/archival storage device via said second local area network without involving said first local area network.

2. An apparatus comprising:

a primary memory;
a mirrored memory;
a backup/archival storage device;
a primary local area network coupled to said primary memory;
a secondary local area network coupled to said

mirrored memory and said backup/archival storage device;

a bridge means for selectively coupling said first local area network to said second local area network

a transaction processor coupled by said primary local area network to said bridge means and said primary memory and programmed to carry out read and write transactions with said primary memory over a primary local area network and, for every write transaction to said primary memory, programmed to carry out a write transaction of the same data to said mirrored memory via said bridge means and said secondary local area network;

a backup and hierarchical storage management processor coupled to said backup/archival storage device and said mirrored memory and said bridge means via said secondary local area network, and programmed to carry out read and write transactions between said mirrored memory and said backup/archival storage device via said second local area network without involving said first local area network.

3. An apparatus comprising:

a primary memory;
a mirrored memory;
a backup/archival storage device;
a primary local area network coupled to said primary memory;
a secondary local area network coupled to said mirrored memory;
a bus coupled to said backup/archival storage device;
a third local area network;
a backup and hierarchical storage management processor coupled to said third local area network, and coupled to said backup/archival storage device via said bus and coupled to said mirrored memory via said secondary local area network, and programmed to carry out read and write transactions between said mirrored memory and said backup/archival storage device via said second local area network and said bus without involving said first local area network;
a transaction processor coupled by said primary local area network to said primary memory and coupled to said backup and hierarchical storage management processor via said third local area network, and programmed to carry out read and write transactions with said primary memory over said primary local area network and, for every write transaction to said primary memory, said transaction processor is programmed to carry out a mirrored write transaction of the same data to said mirrored memory

by sending the same data written on said primary memory to said backup and hierarchical storage management processor via said third local area memory;
 and wherein said backup and hierarchical storage management processor is programmed to receive data transmissions of said mirrored write transactions and write said data to said mirrored memory via said secondary local area network.

4. An apparatus comprising:

a primary memory;
 a mirrored memory;
 a backup/archival storage device;
 a primary local area network coupled to said primary memory;
 a secondary local area network coupled to said mirrored memory;
 a bus coupled to said backup/archival storage device;
 a transaction processor coupled by said primary local area network to said primary memory, and coupled to said backup/archival storage device via said bus, and coupled to said mirrored memory via said secondary local area network, and programmed to carry out read and write transactions with said primary memory over said primary local area network and, for every write transaction to said primary memory, said transaction processor being programmed to carry out a mirrored write transaction of the same data to said mirrored memory by sending the same data written on said primary memory to said mirrored memory via said secondary local area network, and further programmed to carry out read and write transactions of backup and hierarchical storage management functions between said mirrored memory and said backup/archival storage device using said secondary local area network and said bus without involving said primary local area network.

5. An apparatus comprising:

a Transaction Server computer, said Transaction Server computer being a Fibre Channel Arbitrated Loop node;
 a first array of one or more disk drives or other storage media, said first array being a Fibre Channel Arbitrated Loop node;
 a first Fibre Channel Arbitrated Loop coupling said Transaction Server computer to said one or more disk drives;
 a backup server computer, said backup server computer being a Fibre Channel Arbitrated Loop node;

a second array of one or more disk drives or other storage media which mirror the storage capacity of said first array of one or more disk drives or other storage media, said second array being a Fibre Channel Arbitrated Loop node;

a second Fibre Channel Arbitrated Loop coupling said backup server computer to said second array of one or more disk drives or other storage media;

a Bridge compatible with the Fibre Channel Arbitrated Loop protocol which couples said first Fibre Channel Arbitrated Loop to said second Fibre Channel Arbitrated Loop, said Bridge including means for receiving OPN primitives originating from a node on said first Fibre Channel Arbitrated Loop and forwarding said OPN primitive onto said second Fibre Channel Arbitrated Loop if the destination address of said OPN primitive is a node on said second Fibre Channel Arbitrated Loop but not otherwise, and for receiving OPN primitives originating from a node on said second Fibre Channel Arbitrated Loop and forwarding said OPN primitive onto said first Fibre Channel Arbitrated Loop if the destination address of said OPN primitive is a node on said first Fibre Channel Arbitrated Loop but not otherwise.

6. The apparatus of claim 5 wherein said Transaction Server computer is programmed to initiate a mirrored write transaction to said second array of one or more disk drives or other storage media automatically each time said Transaction Server computer initiates a write transaction to said first array of one or more disk drives or other storage media, said mirrored write transaction being the same data that was written to said first array, said mirrored write transaction being initiated by transmitting an OPN primitive on said first Fibre Channel Arbitrated Loop having a destination address which is the node address of said second array of one or more disk drives or other storage media.

7. The apparatus of claim 5 further comprising one or more workstation computers and wherein said Transaction Server computer is coupled to one or more workstation computers by a local area network.

8. The apparatus of claim 5 further comprising one or more backup data storage devices and wherein said backup server is coupled to said one or more backup data storage devices by a local area network or a SCSI bus.

9. The apparatus of claim 8 wherein said backup server computer is programmed to perform data backup

operations by initiating read transactions addressed to said second array of one or more disk drives or other storage media and taking the data received as a result of said read transactions and writing said data to said backup data storage devices. 5

10. The apparatus of claim 8 wherein said backup server computer is programmed to perform hierarchical storage management operations by scanning the file structures of data stored on said one or more disk drives or other storage media of said second array and initiating selective read transactions for predetermined data stored on said one or more disk drives or other storage media of said second array, said selective read transactions addressed to said second array of one or more disk drives or other storage media and taking the data received as a result of said read transactions and writing said data to predetermined one of said backup data storage devices. 10 15 20

11. An opportunity comprising:

first and second Fibre Channel Arbitrated Loops; 25
a first array of disk drives coupled as a mode on said first Fibre Channel Arbitrated Loop;
a second array of disk drives coupled as a mode on said second Fibre Channel Arbitrated Loop; 30
a first transaction processing means for performing read transactions with said first array of disk drives and for performing write transactions with said first and second arrays of disk drives; 35
Bridge means coupled to said first and second Fibre Channel Arbitrated Loops for coupling said first and second Fibre Channel Arbitrated Loops when said first transaction processing means is performing a write transaction with said second array of disk drives but for isolating said first and second Fibre Channel Arbitrated Loops at other times. 40


12. The apparatus of claim 11 further comprising storage management means coupled to said second Fibre Channel Arbitrated Loop for performing hierarchical storage management functions and data backup transactions with said second array of disk drives using only said second Fibre Channel Arbitrated Loop during intervals when said first transaction processing means is not performing write transactions with said second array of disk drives. 45 50


13. A method of performing hierarchical storage management and data backup functions on a Fibre Channel Arbitrated Loop network, comprising the steps of: 55

dividing said Fibre Channel Arbitrated Loop network into first and second Fibre Channel Arbitrated Loop subnetworks coupled by a Bridge, each of said first and second subnetwork having coupled thereto first and second disk array, respectively, said first subnetwork also having coupled thereto a transaction processor and said second subnetwork having coupled thereto a hierarchical storage management and data backup processor;
using said transaction processor to perform write transactions to said first disk array, and, each time a write transaction is performed to said first disk array, using said transaction processor to perform a write transaction of the identical data to said second disk array;
using said hierarchical storage management and data backup processor to perform hierarchical storage management on data backup transactions, or both, with said second array of disk drives.

14. A method of reading and writing data, comprising:

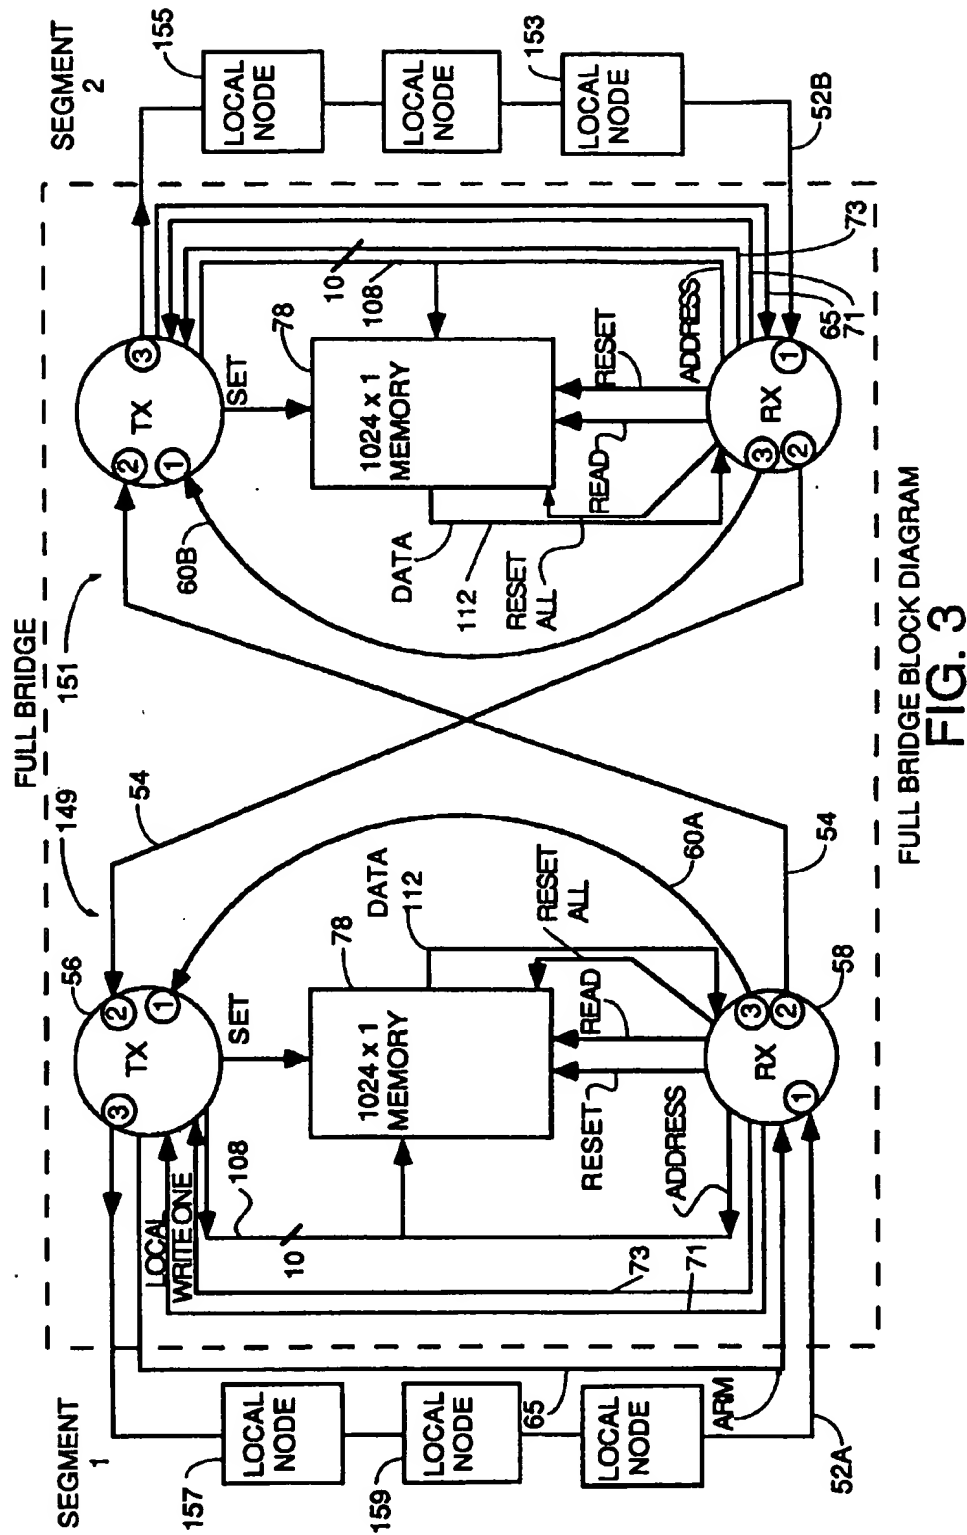
writing data to and reading data from a first memory using a first local area network;
whenever data is written to said first memory using said first local area network, writing the same data to a second memory located on a second local area network using a bridge which selectively couples said first local area network to said second local area network; and
performing backup write transactions between said second memory and a third memory using said second local area network while isolating said second local area network from said first local area network using said bridge.

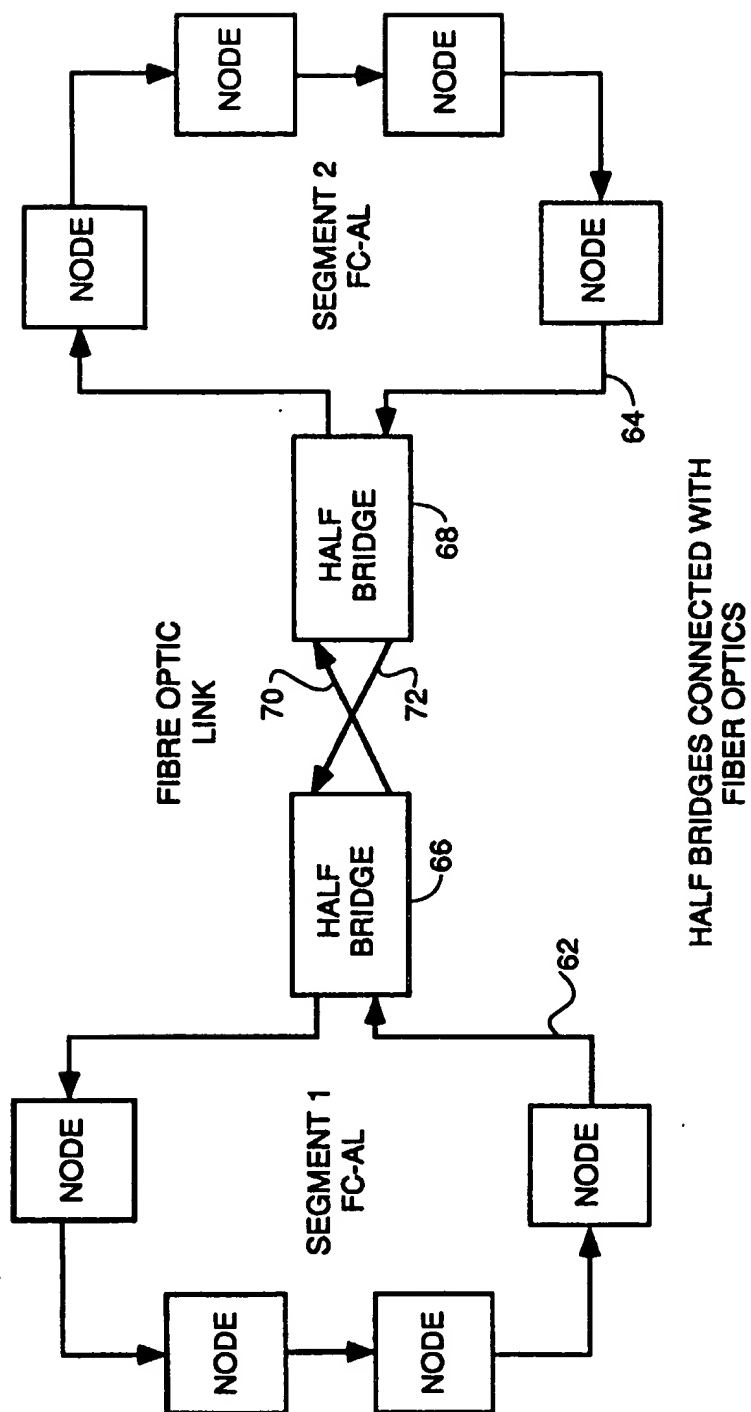




HALF BRIDGE BLOCK DIAGRAM

FIG. 2





HALF BRIDGES CONNECTED WITH
FIBER OPTICS

FIG. 4




Fig. 5

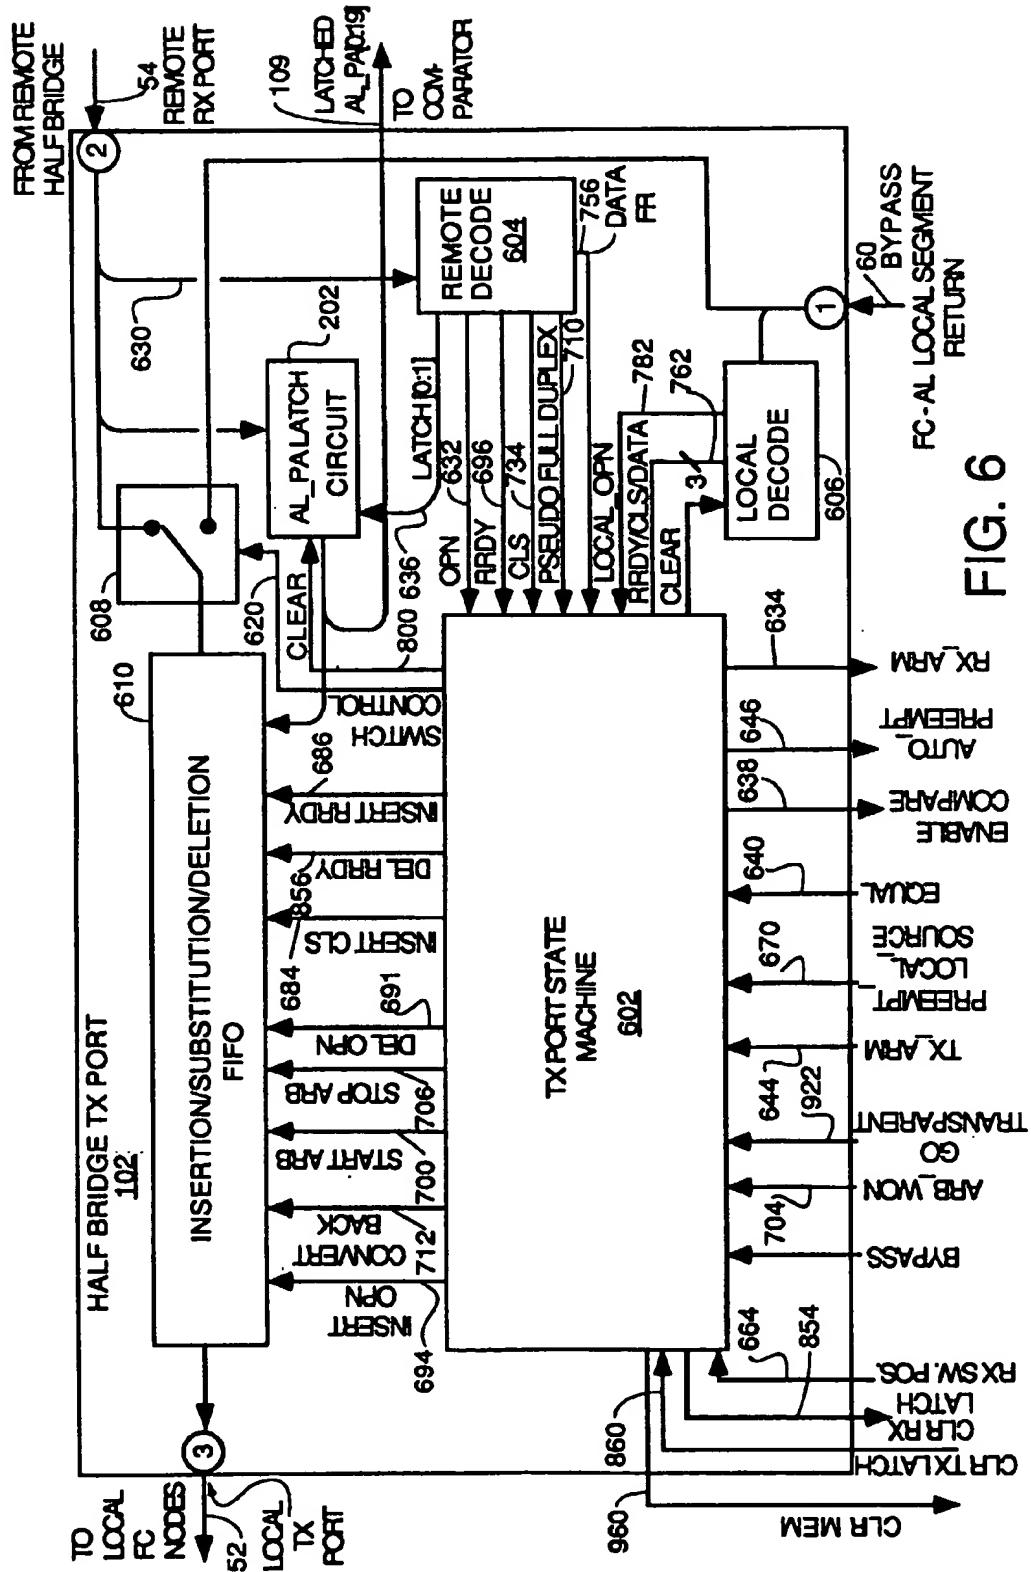


FIG. 6

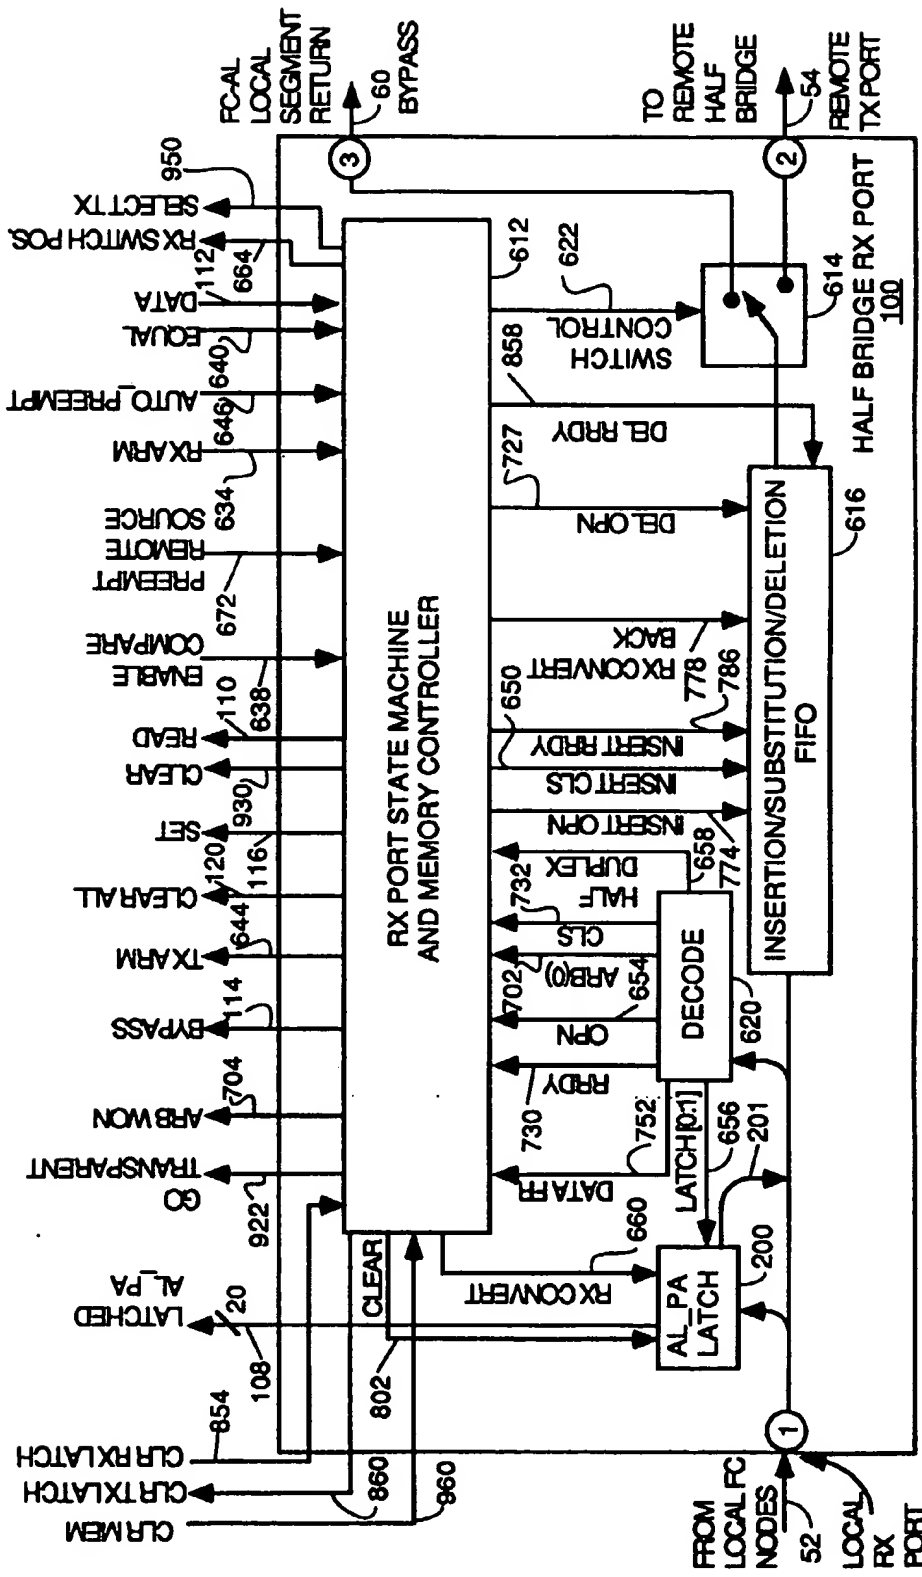


FIG. 7

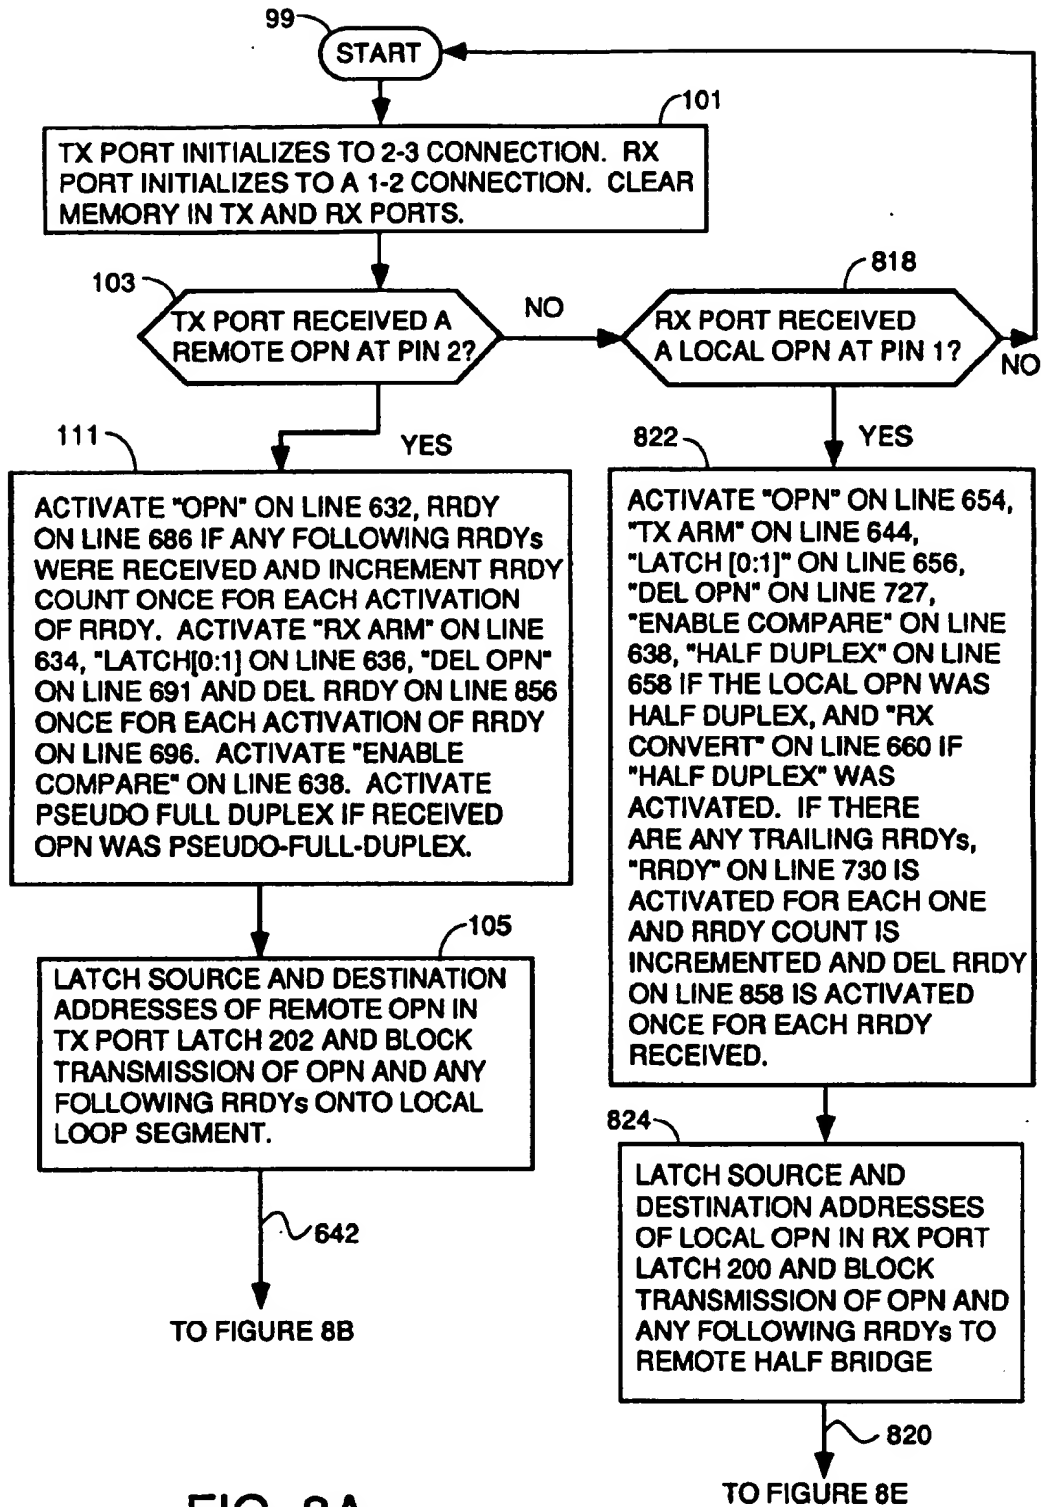
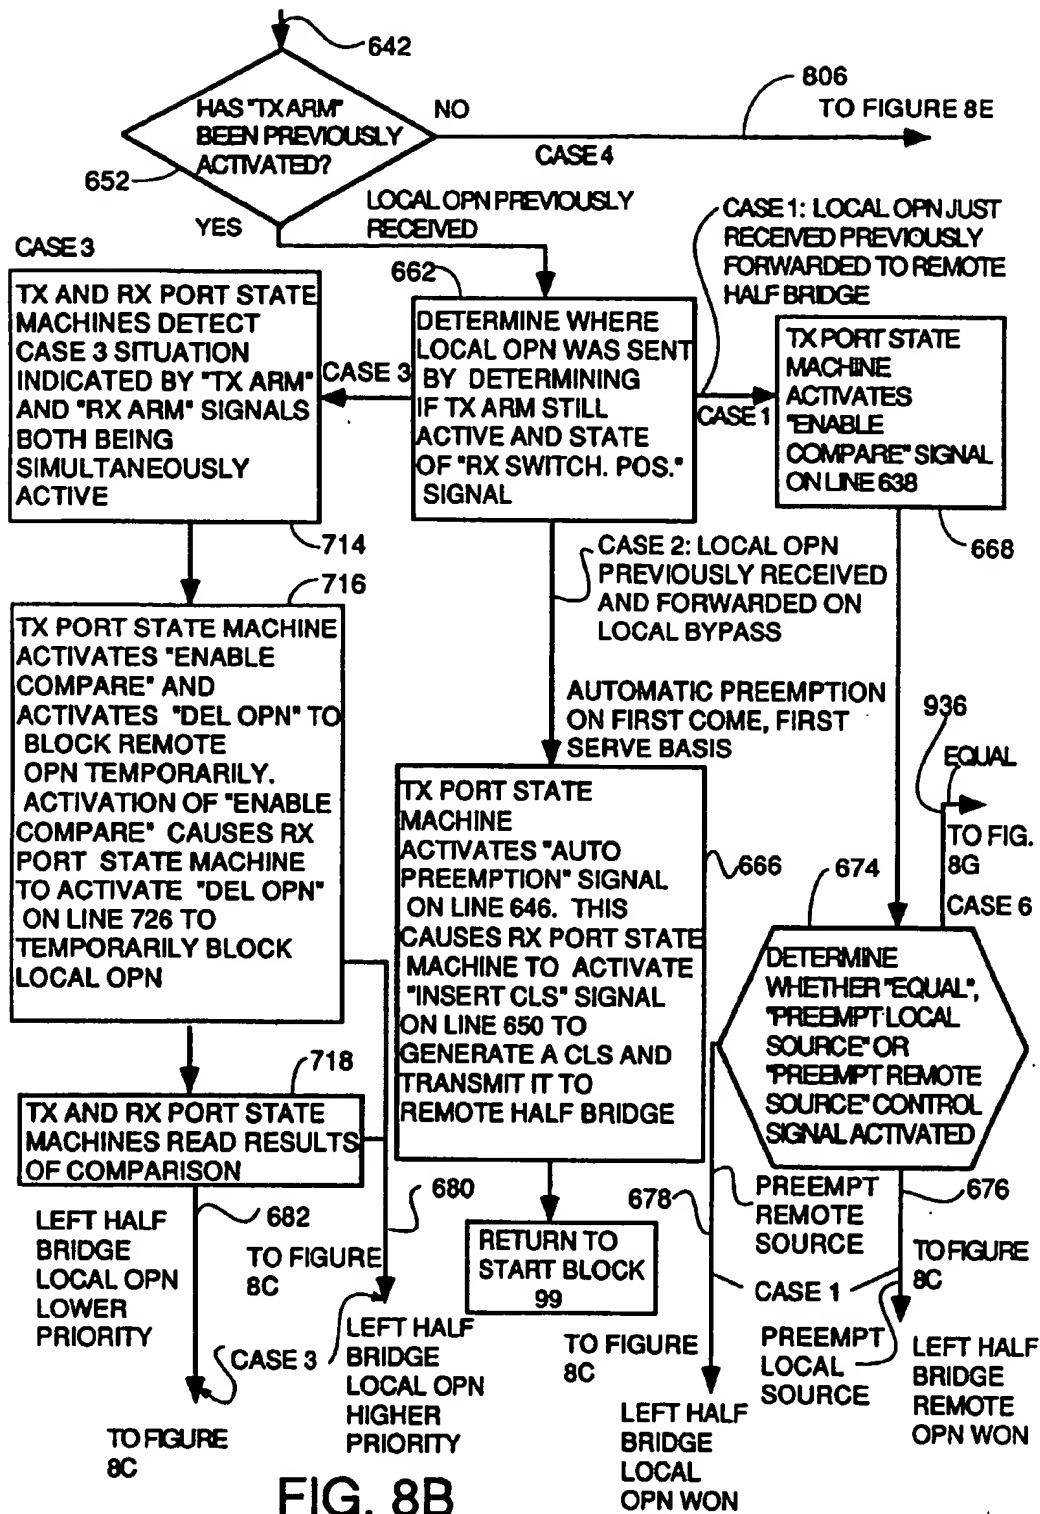
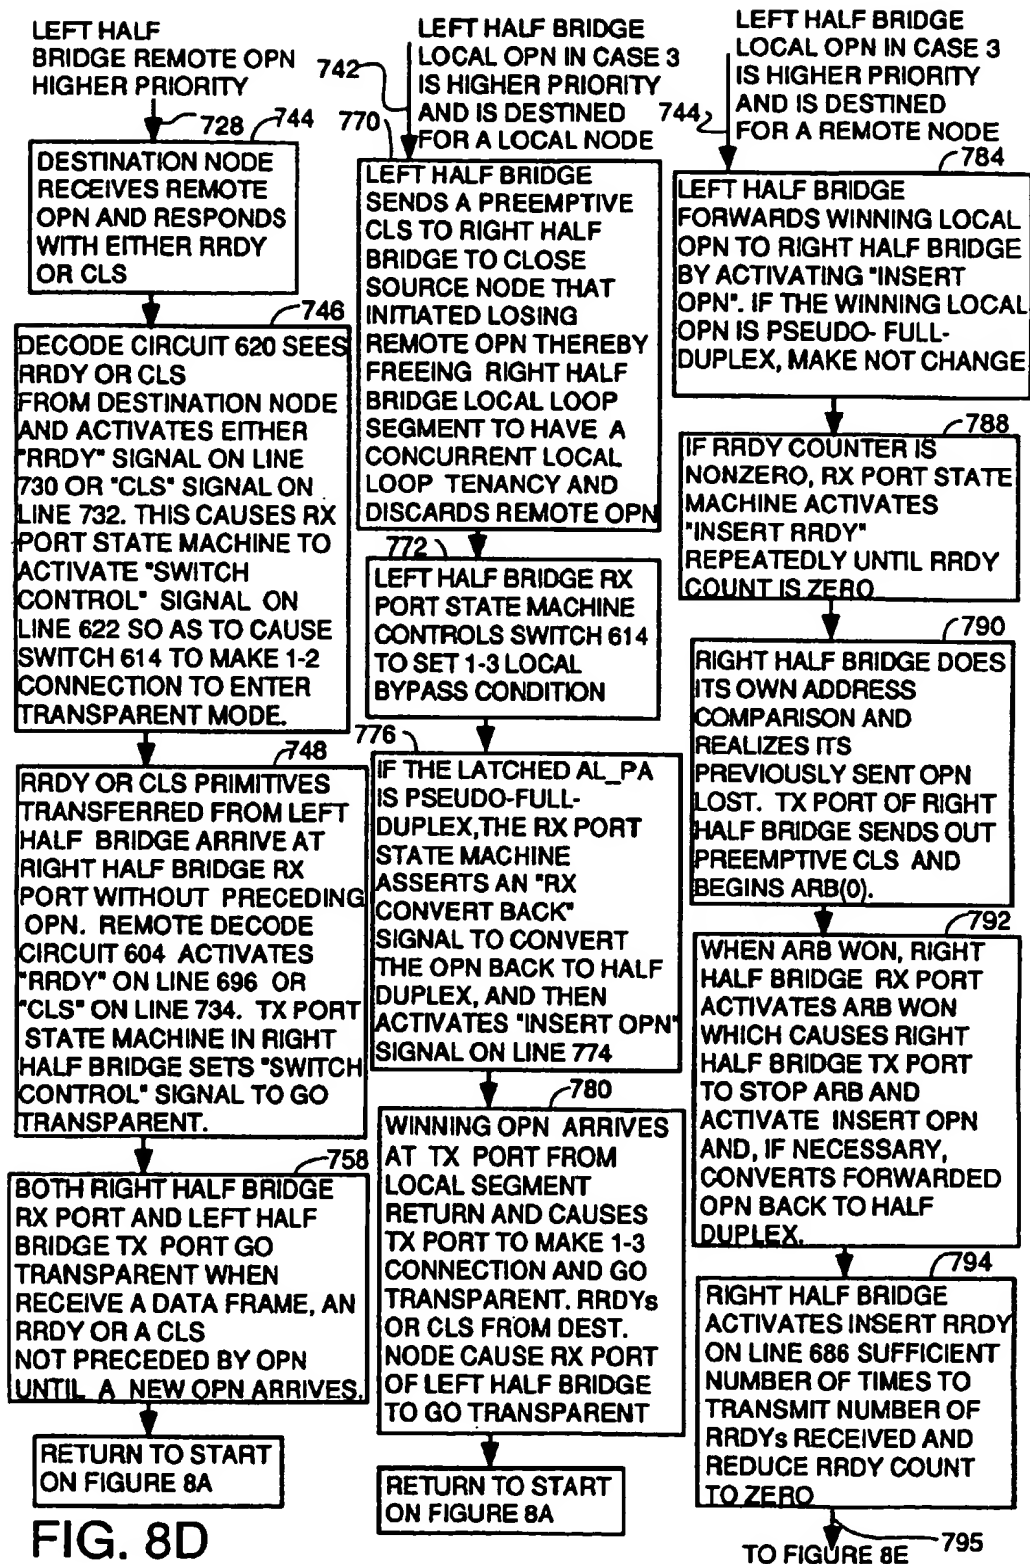
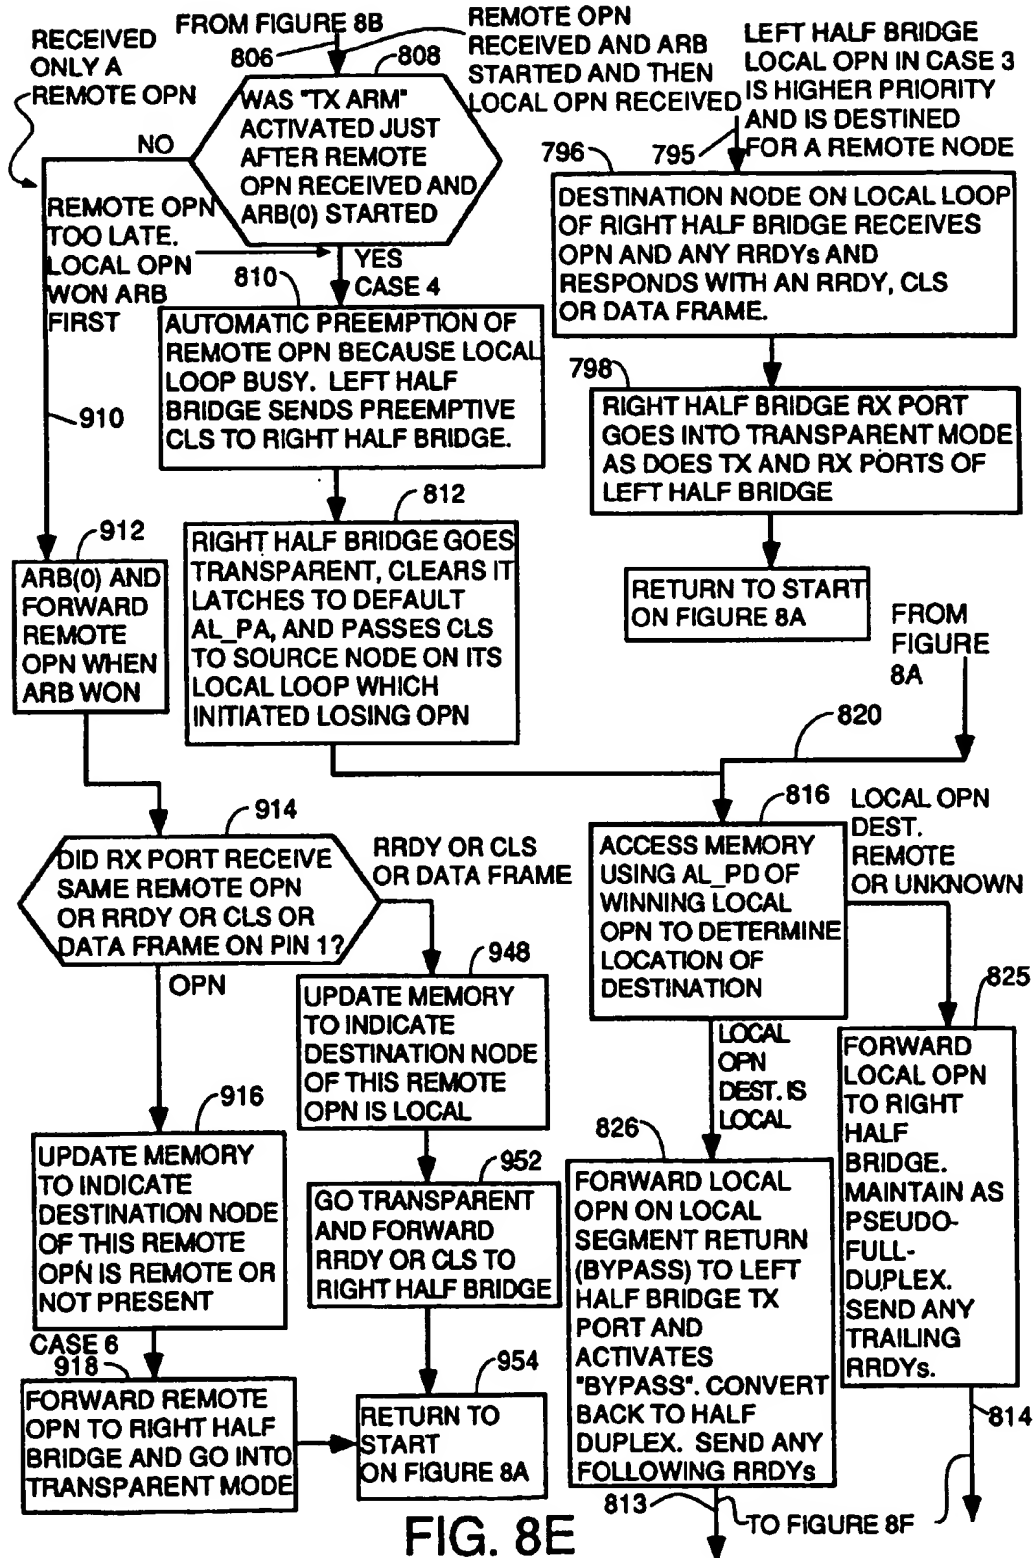


FIG. 8A









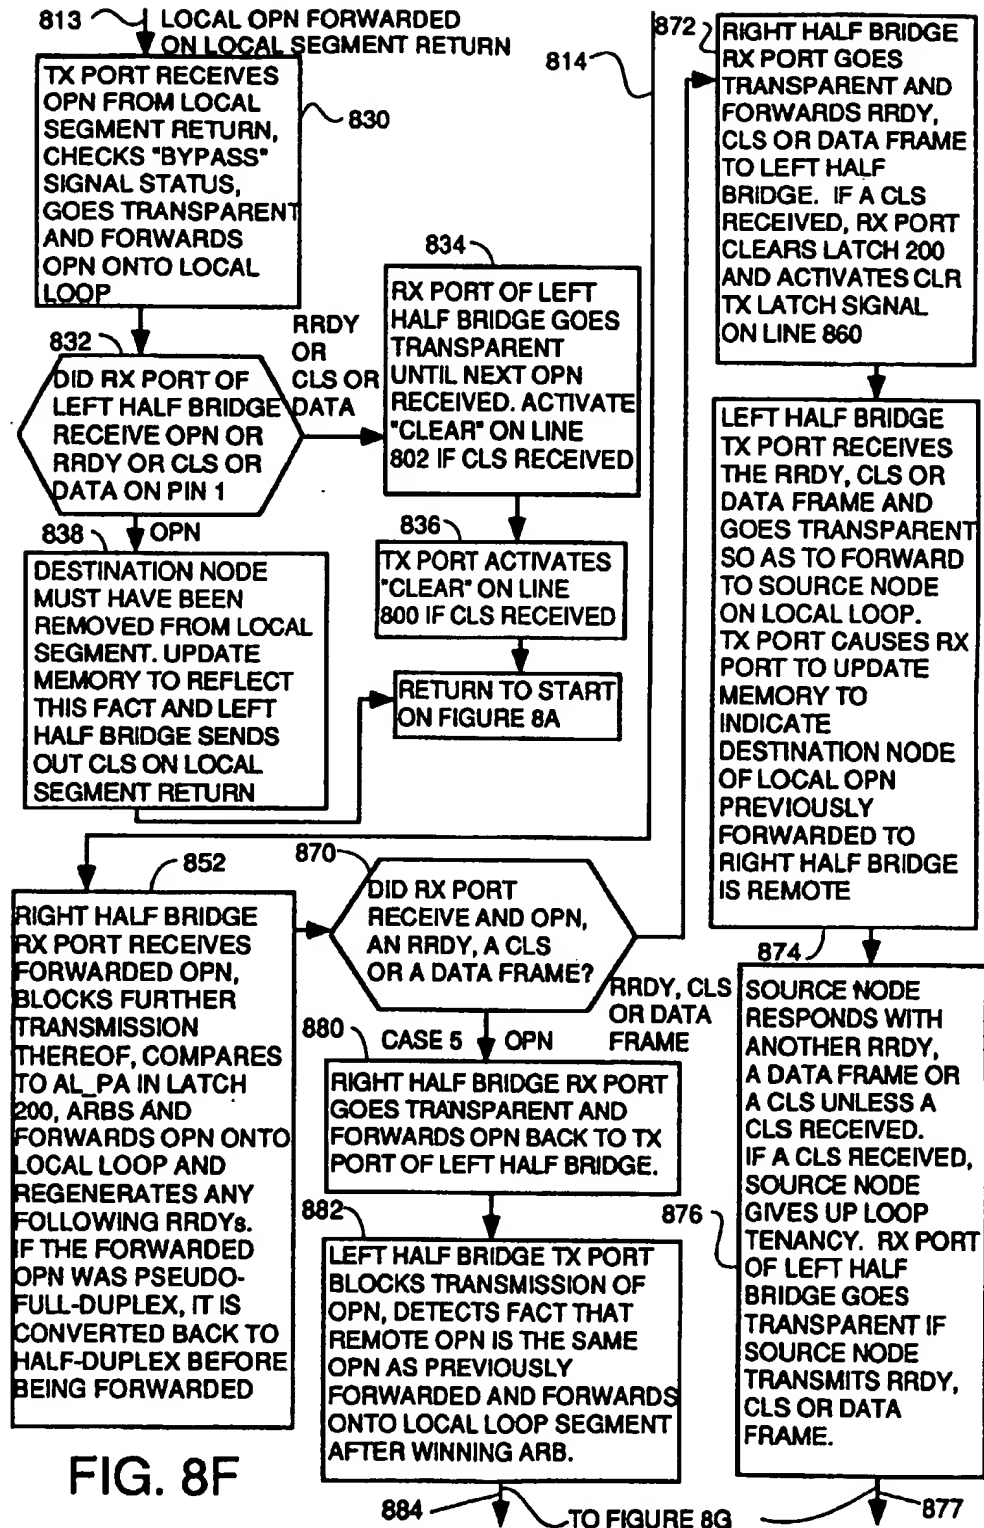
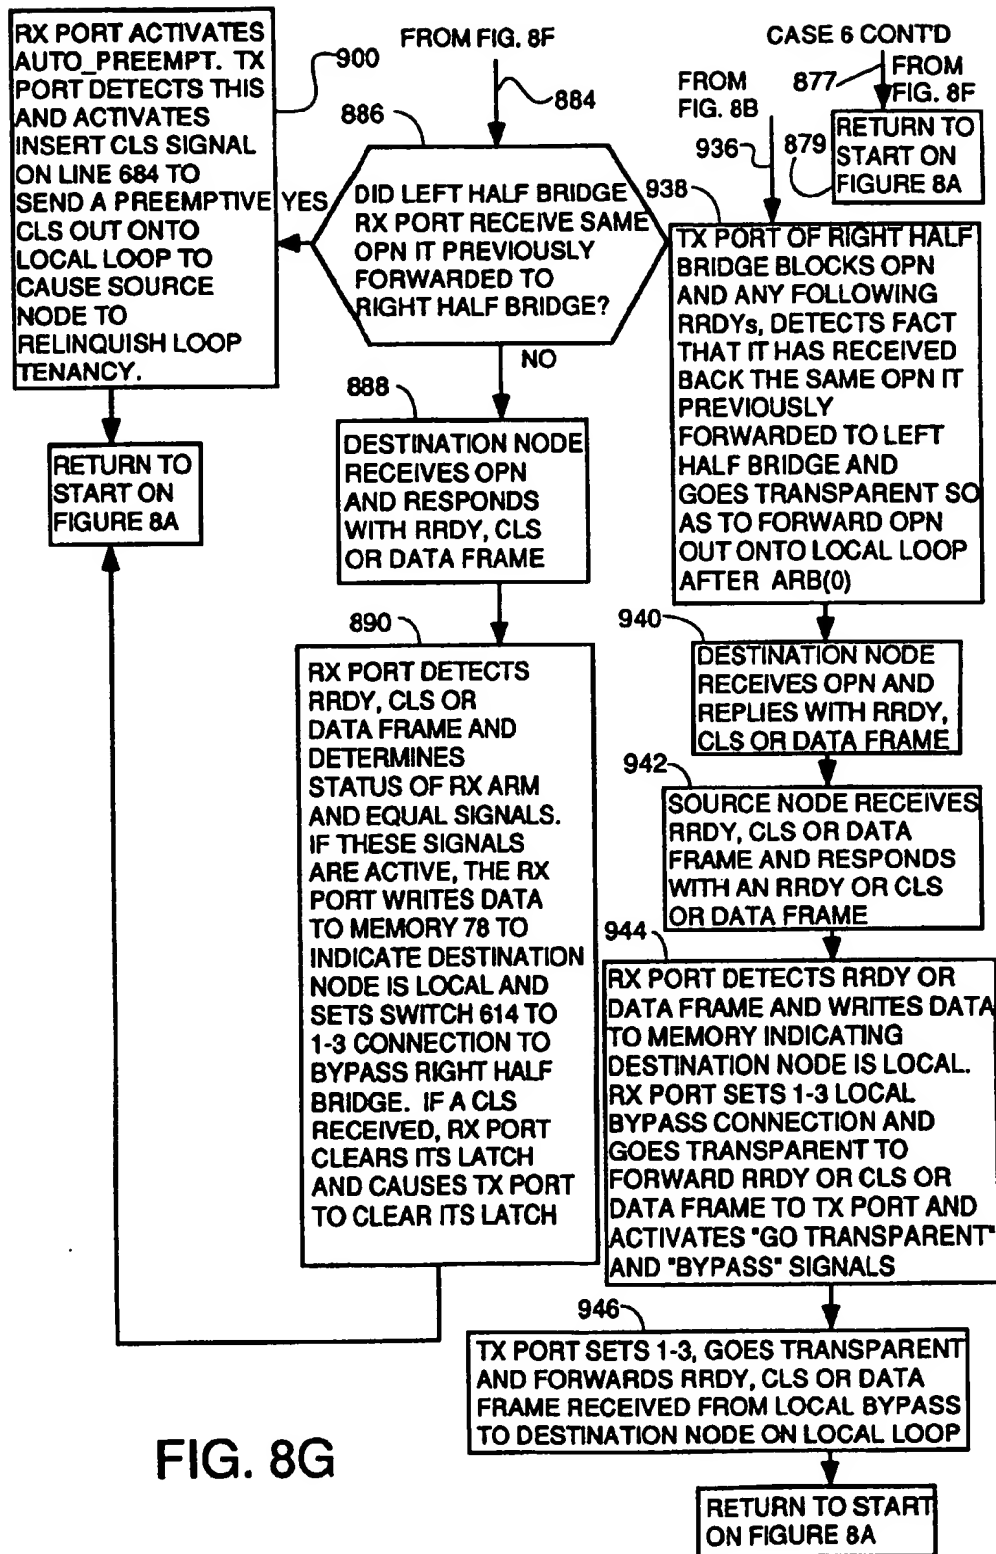


FIG. 8F






Fig. 9

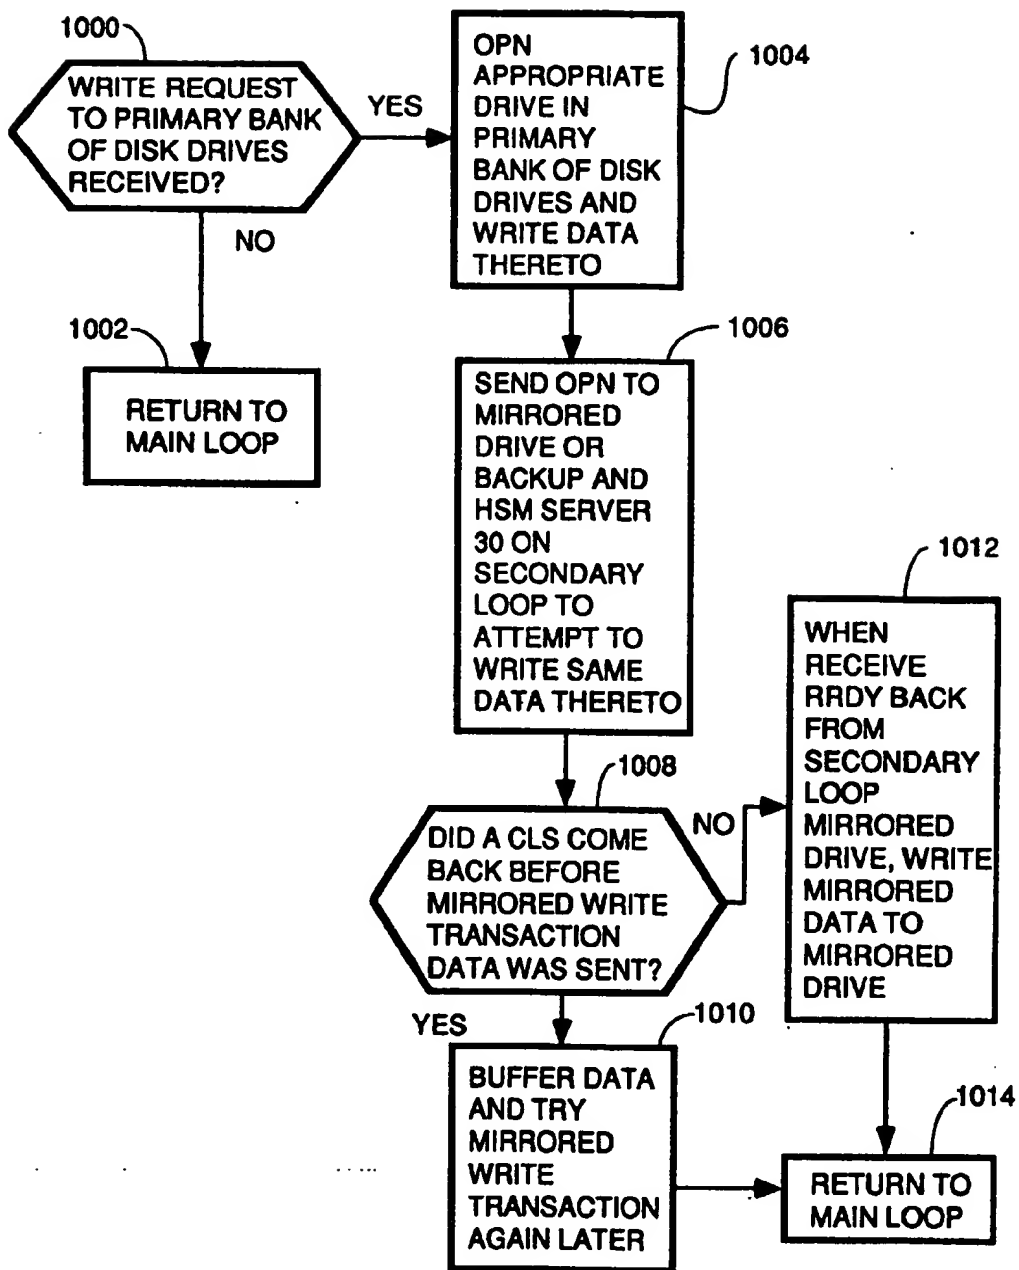


FIG. 10




FIG. 11




FIG. 12

European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 98 66 0029

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
Y A	US 5 212 784 A (SPARKS CLYDE R) 18 May 1993 * abstract * * column 3, line 35 - line 44 * * column 4, line 4 - line 14 * * column 5, line 34 - column 6, line 33 * * figure 2 * ---	1 2-6, 8-10,14	G06F11/14
Y A	EP 0 359 471 A (COMPAQ COMPUTER CORP) 21 March 1990 * abstract * * column 1, line 24 - column 2, line 3 * * figure 1 * ---	1 2-4,14	
A	WO 96 42019 A (OCTOPUS TECHNOLOGIES INC) 27 December 1996 * abstract * * page 5, line 26 - page 6, line 6 * * figure 2 * ---	1-7,14	
A	WO 94 25919 A (EMC CORP) 10 November 1994 * abstract * * page 6, line 26 - page 7, line 15 * * page 8, line 9 - line 22 * * page 12, line 8 - line 20 * * page 13, line 7 - line 13 * * figure 1 * -----	1-6,11, 14	TECHNICAL FIELDS SEARCHED (Int.Cl.6) G06F
The present search report has been drawn up for all claims			
Place of search BERLIN		Date of completion of the search 20 May 1998	Examiner Masche, C
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application I : document cited for other reasons & : member of the same patent family, corresponding document			

81401 (FORM 1503-03-02 (P/M/C/3))